REALIZATION CALCULUS

Komorebi Physics

Continuum, Geometry and Refinement

FIELDS / STRUCTURE / LIMITS

Komorebi P.

2026

KOMOREBI PHYSICS

Komorebi Physics

Continuum, Geometry and Refinement

Komorebi P.

A mathematical account of resolved dynamics, free fields, gauge transport, gravity, arithmetic structure, and cosmological models.

2026

ABOUT THIS VOLUME

Mathematics, models, and evidence

The volume brings together a shared mathematical language and several physical applications. Projector identities and transfer maps lead to complete state spaces, unbounded operators, compatible refinements, and controlled comparisons between descriptions.

Written proofs, recorded compiler checks, finite computations, and physical interpretations have different roles. The mathematical chapters state their hypotheses. The source studies identify the formal declarations, recorded evidence, and remaining verification limits. The visual studies illustrate specified examples rather than replace their proofs.

The free Dirac group, maximal generator domain, and strong Schrödinger equation have compiler-checked proofs. The centered Fourier family also has checked graph-domain and strong-resolvent limits and quadratic bounds under a finite third-order moment. The abstract spectral-functional-calculus identification remains a written deduction without a separate Lean API linkage. Physical-model chapters retain the production-energy obstruction and their stated assumptions.

The equation, theorem, figure, and local section numbers belong to their named treatment. The contents and book folios provide a single navigation system. Source identifiers are retained in the verification material where they are needed for reproducibility.

Copyright © 2026 Komorebi P.
Original text and figures: Creative Commons Attribution 4.0 International.
Original code: MIT License. Third-party material retains its own terms.
creativecommons.org/licenses/by/4.0/

INTRODUCTION

A common language for resolved physics

A physical description often begins by choosing what to resolve. A state may be separated into spatial cells, Fourier modes, spin channels, or arithmetic registers. The useful questions begin after that choice: how is the state reconstructed, how do its components exchange information, and what is retained when the resolution changes?

This book develops those questions through a sequence of concrete constructions. Projector and transfer identities supply the finite algebra. Hilbert-space completion and operator domains supply the analytic control. The free Dirac model carries that route through to a self-adjoint Hamiltonian and the identification of its unitary evolution.

Gauge transport and constrained gravity introduce further requirements. Paths must keep their order, constraints must retain their hypotheses, and a comparison between resolutions must name its topology. Arithmetic and polyphase refinement show how compatible embeddings differ from elimination of hidden degrees of freedom. The latter can create self-energy and memory rather than a smaller autonomous copy of the original dynamics.

The final physical-model chapters ask what follows when a refinement fiber is coupled to spacetime and used in cosmology. Their interpretation is controlled by the production-energy test. A consistent transport calculation from a prescribed distribution does not establish that a microscopic crossing produced that distribution.

The chapters share a language, not a claim that all their physical problems have the same solution. Exact identities, analytic limits, compiler certificates, numerical examples, and proposed interpretations remain distinguishable. The book can therefore be read as one mathematical argument without treating every application as an established theory of nature.

READING THE BOOK

The argument and its conventions

From local structure to complete spaces

Read the first part in order. Packet identities establish what is conserved. Fourier completion makes the state space complete. The operator-domain argument makes the Hamiltonian meaningful. The generator proof then connects that operator to the constructed evolution.

From structure to a physical model

The middle parts change the model while retaining the same discipline: name the state space, the transfer or refinement map, the mathematical assumptions, and the residual obstruction. In the matter and cosmology part, the distribution used for transport remains an explicit input until a controlled production calculation supplies it.

Notation and references

Each named treatment retains its local section, equation, figure, and theorem numbering. References to those numbers are local to that treatment unless another title is named. Preliminary pages use Roman numerals. Arabic page numbers run continuously through the main text and technical appendices. The PDF bookmarks provide detailed navigation within each treatment. The running headers link back to the contents.

Proof and verification

The mathematical chapters carry their hypotheses and written proofs. The visual studies illustrate declared examples. The technical appendices identify the formal statements and recorded computations, including missing evidence and uncompiled extensions. A document build is not a theorem check.

The free-field comparison

The continuum and centered Fourier families use the same infinite mode index set. A changing finite sampled lattice is not silently substituted for that setting. The first-order Hamiltonian domain and the third-order moment required by the quadratic rate are also kept distinct.

CONTENTS

Mathematical constructions

I Resolution and free dynamics

Projectors, Packets, and Transfer 2
Resolution, conservation, and operator reconstruction

Fourier Packets and the Continuum Dirac Operator 35
Clifford matrices, Fourier completion, and the exact domain

The Dirac Group and Its Generator 67
Spectral evolution, resolvents, and regularity-dependent rates

II Gauge transport and gravity

Gauge Covariance and Geometric Refinement 75
Ordered transport, conservation, and curvature estimates

Constrained Gravity and Strong Fields 107
Canonical structure, local reduction, and spherical tests

III Arithmetic and polyphase refinement

Adic and Polyphase Quantum Refinement 161
Haar packets, spectral reduction, and infinite hierarchies

IV Matter and cosmological transport

Covariant Refinement-Fiber Dynamics 213
Symmetry, stress exchange, production, and ultraviolet bounds

Cosmology, Dark Sectors, and Structure Formation 245
Conditional transport, structure, and linked rejection tests

CONTENTS

Visual studies

V Visual studies

Foundations Visual Atlas 396
Packets, currents, reconstruction, and domains

Dirac Continuum Visual Atlas 411
Spinors, Fourier tails, operator domains, and evidence layers

Gauge Visual Atlas 426
Endpoint covariance, path ordering, and curvature

Gravity Visual Atlas 441
Constraints, reduction, and strong-field diagnostics

Adic and Polyphase Visual Atlas 457
Hierarchy, spectral channels, and compatible limits

The visual studies follow the objects developed in the main text. Each example retains its stated inputs and limits. Use the source studies for formal declarations and computational provenance.

CONTENTS

Verification and sources

VI Verification and computational evidence

Verification Scope and Source Conventions 474
Mathematical claims, formal coverage, and physical assumptions

Foundations Formalization Companion 477
Declarations, dependencies, and the foundations certificate

Dirac Formalization Companion 491
The physical Hamiltonian and its exact source correspondence

Gauge Formalization Companion 508
Transport sources, recorded checks, and evidence gaps

Gravity Verification Companion 522
Conditional constructions and computational records

Adic Verification Companion 540
Finite tests and the limits of formal coverage

Fiber Model: Computational Evidence 557
Provenance, numerical methods, claims, and parameters

Cosmology: Reproducibility and Data 564
Source dependencies, reproduction limits, and bibliography

PART I

Resolution and free dynamics

Exact packet algebra becomes field dynamics only after reconstruction, completion, and domain control are supplied. The argument moves from finite transfers to a torus Hamiltonian and then identifies its full unitary group.

IN THIS PART

Projectors, Packets, and Transfer 2

Fourier Packets and the Continuum Dirac Operator 35

The Dirac Group and Its Generator 67

Projectors, Packets, and Transfer

Resolution, conservation, and the passage to infinite systems

Resolution, dynamics, and reconstruction

A physical description often begins by separating a state into spatial cells, frequency bands, spin channels, tensor factors, or spectral sectors. Such a separation is useful only when the components can be reconstructed, their interactions can be described, and different resolutions can be compared. Those requirements lead to the calculus developed here.

An orthogonal projector selects a packet. Placing an operator between two projectors gives a transfer block. A skew-adjoint generator determines an antisymmetric current between packet weights. Grouping the labels produces a coarser description in which internal currents cancel. These statements can be understood in a small matrix example before any infinite-dimensional analysis is needed.

The same questions become more demanding when frames move, states become mixed, or the system exchanges information with an environment. Countable resolutions introduce convergence. Unbounded operators introduce domains. Each extension retains part of the finite algebra and adds a hypothesis that must remain visible where it is used.

The free Dirac construction that follows gives these ideas a concrete analytic realization. Fourier modes resolve a spinor field, finite matrices govern its components, and completion leads to a self-adjoint Hamiltonian. Gauge transport, constrained gravity, and arithmetic refinement then test which parts of the same language survive in different settings.

> **The organizing question**
>
> Which quantities can be reconstructed exactly, which balances survive a change of resolution, and which additional hypotheses are required for an infinite or unbounded system?

The mathematical arguments use ordinary complex linear algebra, operator theory, and Hilbert-space analysis. Formal source correspondences and verification records are collected in the technical appendices. Numerical illustrations explain specified examples rather than replacing universal proofs.

Reading routes and conventions

The first route follows the algebra. Begin with orthogonal resolution and block composition, then read the current identity, coarse aggregation, and moving frames. These sections provide the common notation for the rest of the volume. The worked examples keep the states and operators explicit.

The second route follows physical descriptions. Mixed states, tensor marginals, and Lindblad jumps show which information a packet weight retains and which information still resides in coherence. The positive-effect construction then explains how an overlapping resolution can be represented inside a larger sharp one. Classical probabilities arise only after the relevant normalization and closure assumptions have been supplied.

The third route follows analysis. Countable resolutions require an actual summation statement. An unbounded operator requires its domain. A strong limit acts on each fixed vector, while an operator-norm limit controls all unit vectors at once. The final chapters make these distinctions explicit and explain why the finite pointwise algebra $\mathbb{C}^N$ cannot support a nonzero complex-linear derivation obeying the ordinary Leibniz rule, even though matrix corners support nontrivial commutators.

Convention	Meaning
Inner product	$\langle x, y\rangle$ is conjugate-linear in x and linear in y.
Projector	$P_a = P_a^* = P_a^2$, with orthogonality between distinct packet labels.
Transfer block	$L_{ba} = P_b L P_a$ takes a component in packet a to packet b.
State equation	$\dot{\psi} = K\psi$, with $K^* = -K$. The Hamiltonian convention is $K = -iH$ in units $\hbar = 1$.
Packet weight	$p_a = \|P_a\psi\|^2$. It is a probability for a normalized state and a complete orthogonal resolution. It is not an energy expectation by definition.
Current orientation	$J_{b\leftarrow a}$ is inflow into packet b from packet a, so $\dot{p}_b = \sum_a J_{b\leftarrow a}$ in the finite closed setting.
Mixed state	$\rho \geq 0$, $\operatorname{tr}\rho = 1$ in the finite-dimensional sections.
Coarse cell	A union of disjoint fine packet labels. Its projector is their sum, with the relevant convergence and domain hypotheses in the countable setting.
Unbounded action	An expression such as $H\psi$ includes the condition $\psi \in \mathcal{D}(H)$.
Time and units	K has units of inverse time. We set $\hbar = 1$ when identifying $K = -iH$. All plotted examples state their numerical parameters.

Finite-dimensional statements appear with that hypothesis. A finite projector family can also act on an infinite-dimensional Hilbert space, provided the operators and differentiability assumptions in the statement make sense. These are different uses of the word finite.

The three companions serve different purposes. The monograph gives the mathematical argument. The visual atlas makes specific states, currents, and convergence examples visible. The Lean companion records exact declarations and evidence. A displayed ordinary proof is not silently promoted to a newly compiled theorem, and a sampled curve is not used to establish a universal limit.

Finite resolutions and dynamics

1 Resolving a state into orthogonal packets

A packet begins with a choice of resolution. One may separate a field by spatial cells, a vector by coordinate groups, or a wave by spectral bands. The meaning of the labels changes between applications. The algebra below depends on the relations between the projectors and on the operators acting between their ranges. It does not require a preferred coordinate basis inside each range.

Throughout this section, $\mathcal{H}$ is a complex Hilbert space and I is a finite index set. The inner product is conjugate-linear in its first argument and linear in its second. Every operator initially belongs to $\mathcal{B}(\mathcal{H})$, the bounded operators defined on all of $\mathcal{H}$. The dimension of $\mathcal{H}$ may be infinite. Finiteness refers to the number of packets. Countable resolutions and unbounded generators require the additional hypotheses developed later.

Definition 1.1 (Finite orthogonal atlas)**.** A finite orthogonal atlas is a family $(P_a)_{a\in I}$ satisfying

$$P_a^* = P_a, \qquad P_a P_b = \delta_{ab} P_a, \qquad \sum_{a\in I} P_a = \mathrm{id}. \tag{1}$$

Its packet of a vector ψ is $\psi_a = P_a\psi$. Its packet weight is $p_a = \|\psi_a\|^2$. Zero projectors are allowed, so a label can describe an empty sector.

The word weight avoids assigning physical units before an application supplies them. If $\|\psi\| = 1$, these weights form a probability distribution. If the state carries another normalization, they sum to its squared norm. They are not automatically the energy expectation of a Hamiltonian. For example, a normalized energy eigenvector always has total packet weight one, regardless of its eigenvalue.

Proposition 1.2 (Reconstruction and norm separation)**.** *For every $\psi, \phi \in \mathcal{H}$,*

$$\psi = \sum_a \psi_a, \qquad \langle\psi,\phi\rangle = \sum_a \langle P_a\psi, P_a\phi\rangle, \qquad \|\psi\|^2 = \sum_a \|\psi_a\|^2. \tag{2}$$

Packets with different indices are orthogonal, including packets obtained from different vectors.

Proof. Evaluate $\sum_a P_a = \mathrm{id}$ on ψ. For $a \neq b$, self-adjointness and orthogonality give

$$\langle P_a\psi, P_b\phi\rangle = \langle\psi, P_aP_b\phi\rangle = 0.$$

Expand both entries of $\langle\psi,\phi\rangle$ using the finite reconstruction. All terms with unequal indices vanish. Taking $\phi = \psi$ proves the norm identity. □

This proof separates three roles. Completeness reconstructs the vector, self-adjointness relates the projections to the inner product, and pairwise orthogonality removes cross terms. A mere decomposition of the identity by arbitrary linear maps does not imply the norm formula. A family of positive effects supplies a different probability construction, discussed later, because effects need not be idempotent or pairwise orthogonal.

Example 1.3 (Two packets inside four coordinates)**.** Let $\mathcal{H} = \mathbb{C}^4$ with its usual inner product, and set

$$\psi = \frac{1}{\sqrt{7}}(1, i, 2, -i)^{\mathsf{T}}, \qquad P_A = \mathrm{diag}(1,1,0,0), \quad P_B = \mathrm{diag}(0,0,1,1).$$

Then $\psi_A = (1, i, 0, 0)^{\mathsf{T}}/\sqrt{7}$ and $\psi_B = (0, 0, 2, -i)^{\mathsf{T}}/\sqrt{7}$. The weights are $p_A = 2/7$ and $p_B = 5/7$. Refining into the four coordinate projectors gives weights $1/7, 1/7, 4/7, 1/7$. The two descriptions resolve the same vector at different levels. The coarse description still contains the two-dimensional packet vectors. Retaining only the two scalar weights would discard their internal phases and amplitudes.

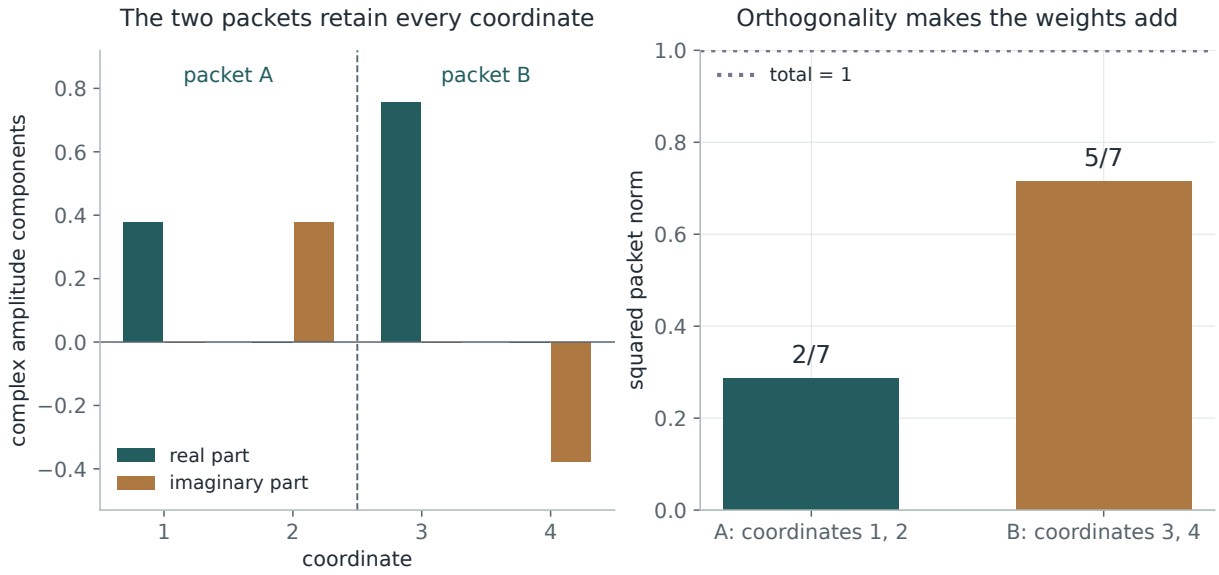


Figure 1. The four-coordinate state of Example 1.3. Packet vectors retain amplitudes and phases, while their two weights retain only the squared sizes.

2 Operator blocks and exact reconstruction

The same atlas resolves an operator on both its input and output sides. For $X \in \mathcal{B}(\mathcal{H})$, define

$$X_{ba} = P_b X P_a. \tag{3}$$

The first index names the destination and the second names the source. As a map between packet spaces, this is $X_{ba} : \operatorname{Ran} P_a \to \operatorname{Ran} P_b$. As a bounded operator on $\mathcal{H}$, it is extended by zero on the orthogonal complement of $\operatorname{Ran} P_a$. Keeping these two descriptions aligned prevents a block from being composed with an incompatible input space.

Proposition 2.1 (Block calculus). *For bounded X, Y,*

$$X = \sum_{b,a} X_{ba}, \qquad P_b X \psi = \sum_a X_{ba} \psi_a, \tag{4}$$

$$(XY)_{ca} = \sum_b X_{cb} Y_{ba}, \qquad (X^*)_{ab} = (X_{ba})^*. \tag{5}$$

The identity has diagonal blocks P_a and zero off-diagonal blocks. If $X = X^*$, its opposite blocks are adjoints.

Proof. Insert $\sum_b P_b = \mathrm{id}$ to the left of X and $\sum_a P_a = \mathrm{id}$ to its right. This gives the first formula. Evaluate the right resolution on ψ for the second. For the product,

$$P_c XY P_a = P_c X\Big(\sum_b P_b\Big) Y P_a = \sum_b (P_c X P_b)(P_b Y P_a),$$

where $P_b^2 = P_b$ supplies the repeated middle projector. Finally, $(P_b X P_a)^* = P_a X^* P_b$. □

The multiplication law is ordinary matrix multiplication with operator-valued entries. It records the intermediate packet through which a composition passes. This matters even when two different paths reach the same destination. Their operator amplitudes add before any norm or probability is taken. Replacing each block by its norm would generally destroy interference and would no longer reconstruct the original operator.

Example 2.2 (A coarse block retains internal channels)**.** Use the two-packet split of Example 1.3 and the Hermitian matrix

$$H = \begin{pmatrix} 0 & 1 & 0 & 2 \\ 1 & 0 & -1 & 0 \\ 0 & -1 & 0 & 1 \\ 2 & 0 & 1 & 0 \end{pmatrix}.$$

Relative to $\mathbb{C}^4 = \mathbb{C}^2_A \oplus \mathbb{C}^2_B$, the four blocks are

$$H_{AA} = H_{BB} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad H_{BA} = \begin{pmatrix} 0 & -1 \\ 2 & 0 \end{pmatrix}, \qquad H_{AB} = H_{BA}^*.$$

The block from A to B contains two distinct coordinate channels. Neither has to vanish for their net probability current to vanish at a particular state. The operator block describes available coupling. The current will also depend on the state occupying its source and destination.

Deleting off-diagonal blocks produces another operator, $\sum_a P_a X P_a$. It agrees with X exactly when all omitted blocks vanish. Thus an atlas resolution is lossless, while keeping only its diagonal is an additional operation. Similarly, a finite-dimensional coordinate choice inside the packet ranges turns the blocks into numerical submatrices, but that coordinate representation is not needed for reconstruction itself.

3 From skew-adjoint evolution to transfer current

Let $\psi(t)$ be a differentiable trajectory and suppose that

$$\dot{\psi}(t) = K(t)\psi(t), \qquad K(t)^* = -K(t). \tag{6}$$

Each $K(t)$ is bounded. The argument uses the stated derivative of the trajectory at the time under consideration. It does not manufacture a solution of an arbitrary differential equation. For a fixed bounded generator, the exponential series provides a separate construction of the evolution. A state-dependent vector field requires its own existence and uniqueness hypotheses.

Define the receiving-first current by

$$J_{b\leftarrow a}(t) = 2\,\mathrm{Re}\,\langle \psi_b(t), K(t)\psi_a(t)\rangle = 2\,\mathrm{Re}\,\langle \psi_b(t), K_{ba}(t)\psi_a(t)\rangle\,. \tag{7}$$

The second equality follows because $P_b\psi_b = \psi_b$. A positive value means that the contribution assigned to the pair increases the weight in b. The current is a real number, and its sign convention is fixed by the derivative calculation below.

Theorem 3.1 (Finite packet transfer)**.** *For the trajectory in* (6)*,*

$$\dot{p}_b = \sum_a J_{b\leftarrow a}, \tag{8}$$

$$J_{b\leftarrow a} = -J_{a\leftarrow b}, \qquad J_{a\leftarrow a} = 0, \tag{9}$$

$$\frac{d}{dt}\sum_a p_a = 0. \tag{10}$$

For every subset $C \subseteq I$, its total weight obeys

$$\frac{d}{dt}\sum_{b\in C} p_b = \sum_{b\in C}\sum_{a\notin C} J_{b\leftarrow a}. \tag{11}$$

Proof. The projectors are fixed in time, so the product rule for the inner product gives

$$\begin{aligned}\dot{p}_b &= 2\,\mathrm{Re}\left\langle P_b\psi, P_b\dot{\psi}\right\rangle = 2\,\mathrm{Re}\,\langle\psi_b, K\psi\rangle \\ &= \sum_a 2\,\mathrm{Re}\,\langle\psi_b, K\psi_a\rangle\,.\end{aligned}$$

The reconstruction of ψ justifies the last finite sum. Skew-adjointness gives

$$\langle\psi_a, K\psi_b\rangle = -\langle K\psi_a, \psi_b\rangle = -\overline{\langle\psi_b, K\psi_a\rangle}.$$

Taking real parts proves antisymmetry. Putting $a = b$ gives zero diagonal current. Every internal ordered pair in $C\times C$ cancels with its reverse. Splitting the sum over all sources into C and its complement proves (11). Taking $C = I$ proves total conservation. □

The first line of this proof does not require skew-adjointness. It expresses the derivative of the actual packet weight for any bounded generator. Skew-adjointness is the additional input that turns the resulting pairing into an antisymmetric transfer law. Without it, one generally has sources or sinks as well as exchange.

In the Hamiltonian convention $K = -iH$ with $H = H^*$ and $\hbar = 1$, linearity in the second slot yields

$$J_{b\leftarrow a} = 2\,\mathrm{Im}\,\langle\psi_b, H\psi_a\rangle\,. \tag{12}$$

Indeed, $\mathrm{Re}(-iz) = \mathrm{Im}\,z$. The plus sign belongs to this inner-product convention and to the choice that the first index receives the current. It should not be changed independently of either convention.

Example 3.2 (An exactly solved exchange)**.** On $\mathbb{C}^2$, take the coordinate atlas,

$$H = g\begin{pmatrix}0 & 1\\ 1 & 0\end{pmatrix}, \quad g > 0, \qquad \psi(0) = \begin{pmatrix}1\\ 0\end{pmatrix}.$$

Direct differentiation verifies the solution

$$\psi(t) = \begin{pmatrix} \cos(gt) \\ -i\sin(gt) \end{pmatrix}, \quad p_1 = \cos^2(gt), \quad p_2 = \sin^2(gt).$$

The current is

$$J_{2\leftarrow 1} = 2\,\mathrm{Im}(ig\sin(gt)\cos(gt)) = g\sin(2gt) = \dot{p}_2.$$

It reverses sign during the cycle. At $t = 0$ the coupling is nonzero but the instantaneous current is zero, since the destination amplitude vanishes. The population begins to change at second order in time. This is why coupling strength and instantaneous flow must be kept separate.

The same proof applies to $\dot{\psi} = K(\psi, t)\psi$ when the operator evaluated along the supplied trajectory is skew-adjoint. No derivative of the function $K(\psi, t)$ appears in the first derivative of $\|P_b\psi\|^2$. Such derivatives would enter a second derivative or a linearized perturbation equation. Pointwise conservation therefore survives many nonlinear choices of generator, while well-posedness remains a separate question.

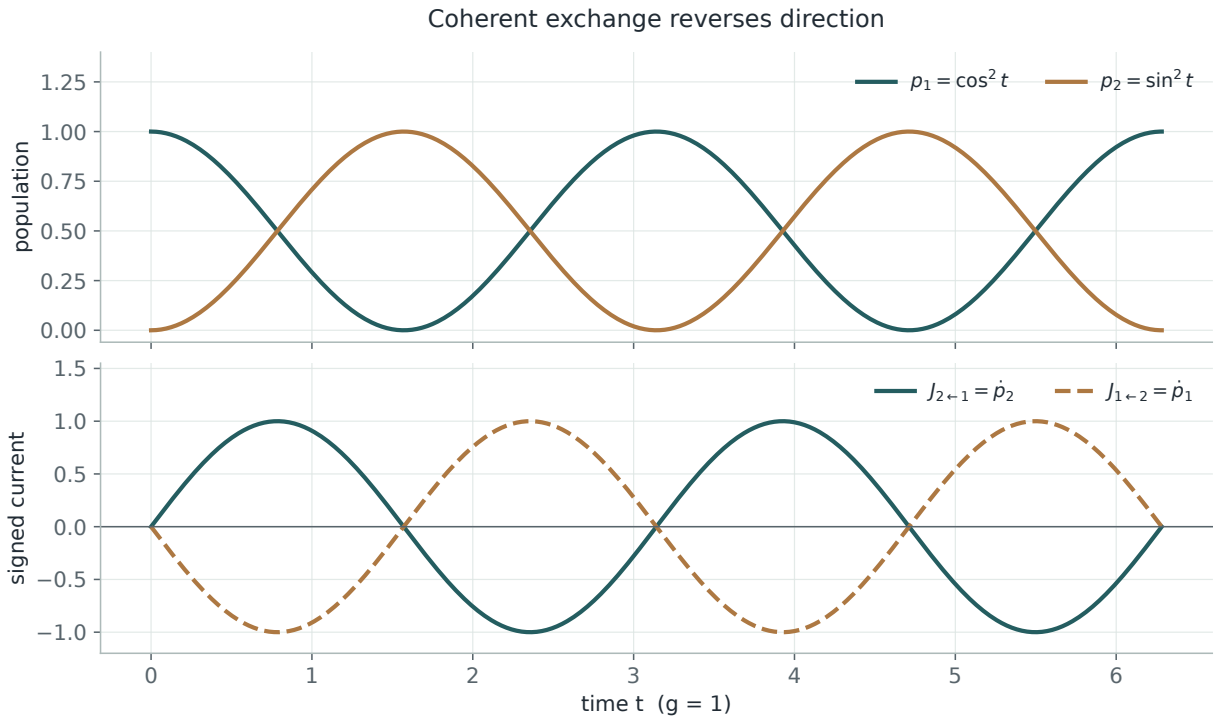


Figure 2. For the two-level Hamiltonian with coupling one, coherent evolution exchanges population periodically. The destination-first current changes sign when the direction of transfer reverses.

4 Coarsening preserves the balance exactly

Let $\mathcal{A}$ be a finite set of coarse labels and let $f : I \to \mathcal{A}$ assign every fine packet to exactly one coarse label. Write $C_A = f^{-1}(A)$ and define

$$Q_A = \sum_{a \in C_A} P_a. \tag{13}$$

The fibers are disjoint and cover I, so (Q_A) is another orthogonal atlas. Empty fibers give zero projectors. This is a grouping of an existing resolution. It need not be an approximation to a different physical system.

Theorem 4.1 (Exact coarse transfer)**.** *The coarse weight and current are*

$$p^{\mathrm{c}}_A = \|Q_A\psi\|^2 = \sum_{a\in C_A} p_a, \tag{14}$$

$$J^{\mathrm{c}}_{B\leftarrow A} = 2\,\mathrm{Re}\,\langle Q_B\psi, KQ_A\psi\rangle = \sum_{b\in C_B}\sum_{a\in C_A} J_{b\leftarrow a}. \tag{15}$$

They satisfy $\dot p^{\mathrm{c}}_B = \sum_A J^{\mathrm{c}}_{B\leftarrow A}$ and $J^{\mathrm{c}}_{B\leftarrow A} = -J^{\mathrm{c}}_{A\leftarrow B}$. In particular, the sum of fine currents internal to a coarse cell is zero.

Proof. Orthogonality proves the weight identity. Finite additivity of the real part of the inner product proves the current identity. The coarse projectors satisfy (1), so Theorem 3.1 applies to them. Equivalently, sum the fine balances inside each fiber and cancel all internal pairs. □

For an incidence matrix $R_{Aa} = 1$ when $f(a) = A$ and zero otherwise, these formulas read

$$p^{\mathrm{c}} = Rp, \qquad J^{\mathrm{c}} = RJR^{\mathsf{T}}, \qquad R^{\mathsf{T}}\mathbf{1}_{\mathrm{c}} = \mathbf{1}_{\mathrm{f}}. \tag{16}$$

The last identity expresses exactly one coarse assignment per fine index. An arbitrary binary matrix does not suffice. If S performs a second grouping, then

$$S(Rp) = (SR)p, \qquad S(RJR^{\mathsf{T}})S^{\mathsf{T}} = (SR)J(SR)^{\mathsf{T}}.$$

Thus a tower of exact groupings is associative. There is no accumulation of approximation error in these identities.

Example 4.2 (Internal exchange disappears from the coarse boundary)**.** Consider four labels with nonzero currents

$$J_{2\leftarrow 1} = 3, \quad J_{3\leftarrow 2} = 2, \quad J_{4\leftarrow 3} = 1, \quad J_{1\leftarrow 4} = \tfrac{1}{2},$$

their negatives in the reverse directions, and all other currents zero. This is an example of the balance algebra, without an asserted underlying quantum trajectory. The weight derivatives are $(-5/2, 1, 1, 1/2)$. For $A = \{1, 2\}$ and $B = \{3, 4\}$, the internal currents 3 and 1 cancel, leaving $J^{\mathrm{c}}_{B\leftarrow A} = 2 - 1/2 = 3/2$. Hence $\dot p^{\mathrm{c}}_A = -3/2$ and $\dot p^{\mathrm{c}}_B = 3/2$.

Exact balance should not be confused with closure of a reduced equation. The scalar vector p^{c} generally does not determine the state-dependent current J^{c}. Internal phase or correlation data may still be needed. Even at the fine level, a continuity equation alone does not uniquely determine its current. On three labels, the antisymmetric matrix

$$\begin{pmatrix} 0 & c & -c \\ -c & 0 & c \\ c & -c & 0 \end{pmatrix}$$

has zero row sums for any real c. Adding it changes circulation while preserving all weight derivatives. The operator and state specify a particular current through (7). Observing only the derivatives cannot recover that additional information.

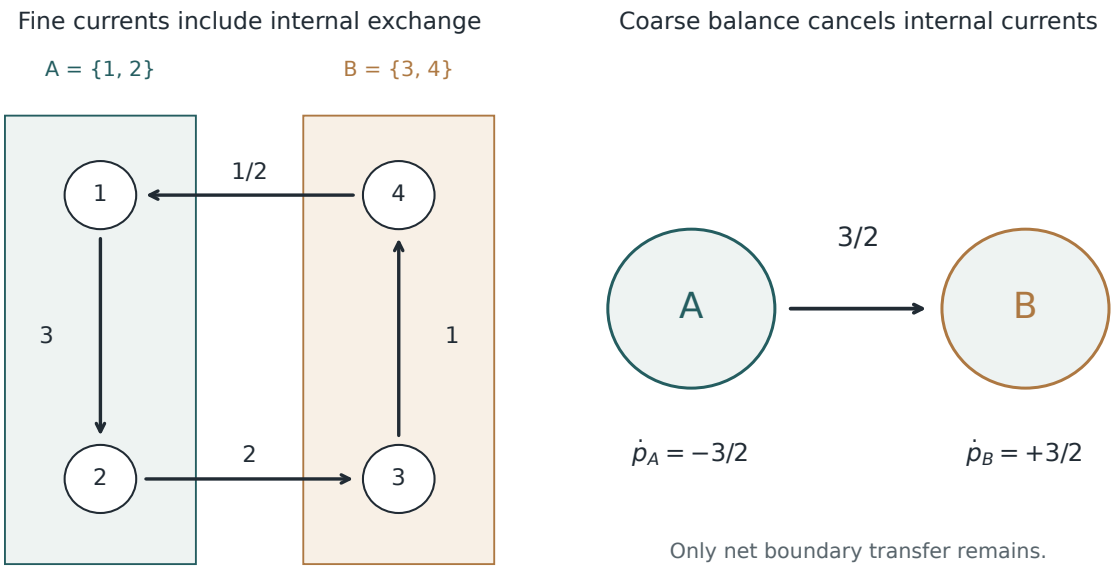


Figure 3. An instantaneous antisymmetric current array and its two-cell aggregation. The internal arrows cancel, leaving net transfer $3/2$ across the coarse boundary. This balance example does not assume a constructed Hamiltonian trajectory.

5 Moving atlases and their connection

A resolution can change with time even when the state is stationary. To retain orthogonality, transport the entire atlas by a unitary family. Let $U(t)$ be differentiable in operator norm and satisfy both $U^*U = \mathrm{id}$ and $UU^* = \mathrm{id}$. Set

$$P_a(t) = U(t)P_a^{(0)}U(t)^*, \qquad \phi(t) = U(t)^*\psi(t), \qquad \Gamma(t) = U(t)^*\dot{U}(t). \tag{17}$$

The coordinates ϕ see fixed projectors $P_a^{(0)}$. The physical weight in the moving packet is the same scalar:

$$\|P_a(t)\psi(t)\|^2 = \left\|P_a^{(0)}\phi(t)\right\|^2.$$

Proposition 5.1 (The transformed generator)**.** *If $\dot{\psi} = K\psi$ and $K^* = -K$, then*

$$\dot{\phi} = K_{\mathrm{eff}}\phi, \qquad K_{\mathrm{eff}} = U^*KU - \Gamma, \qquad \Gamma^* = -\Gamma, \quad K_{\mathrm{eff}}^* = -K_{\mathrm{eff}}. \tag{18}$$

Consequently the fixed-atlas current formula, applied to ϕ and K_{eff}, gives the derivative of the moving physical packet weights.

Proof. Differentiate $U^*U = \mathrm{id}$ to obtain $\dot{U}^*U + U^*\dot{U} = 0$. This proves $\Gamma^* = -\Gamma$. Multiplication on the right by U^* gives $\dot{U}^* = -\Gamma U^*$. Therefore

$$\dot{\phi} = \dot{U}^*\psi + U^*\dot{\psi} = -\Gamma\phi + U^*KU\phi.$$

Unitary conjugation preserves skew-adjointness. The difference of the two skew-adjoint operators in (18) is skew-adjoint, so the transfer theorem applies. □

The connection Γ records the velocity of the frame. Omitting it would interpret a change of packet coordinates as if it were entirely a change of the physical state. In Hamiltonian notation the corresponding effective Hamiltonian is $H_{\text{eff}} = U^* H U - i\Gamma$, which is self-adjoint because Γ is skew-adjoint.

Example 5.2 (A stationary vector viewed through a rotating split)**.** Let $\psi(t) = e_1$, $K = 0$, and

$$U(t) = \begin{pmatrix} \cos\theta(t) & -\sin\theta(t) \\ \sin\theta(t) & \cos\theta(t) \end{pmatrix}.$$

With the coordinate projectors as $P_a^{(0)}$,

$$\phi = \begin{pmatrix} \cos\theta \\ -\sin\theta \end{pmatrix}, \quad \Gamma = \dot\theta \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \quad K_{\text{eff}} = -\Gamma.$$

The moving weights are $p_1 = \cos^2\theta$ and $p_2 = \sin^2\theta$, with $J_{2\leftarrow 1} = \dot\theta \sin(2\theta) = \dot p_2$. The state has not moved in the original Hilbert space. The measured sectors have moved around it. The connection term accounts for the entire change.

The same principle governs a time-dependent operator block. Write $V = \dot U U^*$ for the frame velocity in laboratory coordinates. Then

$$\dot P_a = [V, P_a], \qquad \frac{d}{dt}(P_b X P_a) = [V, P_b] X P_a + P_b \dot X P_a + P_b X [V, P_a].$$

Both packet boundaries contribute, as does the operator itself. Conjugation by a general invertible family also preserves idempotence and completeness, but it need not preserve self-adjointness or orthogonality. The unitary hypothesis is what keeps these transported objects within the present orthogonal atlas setting.

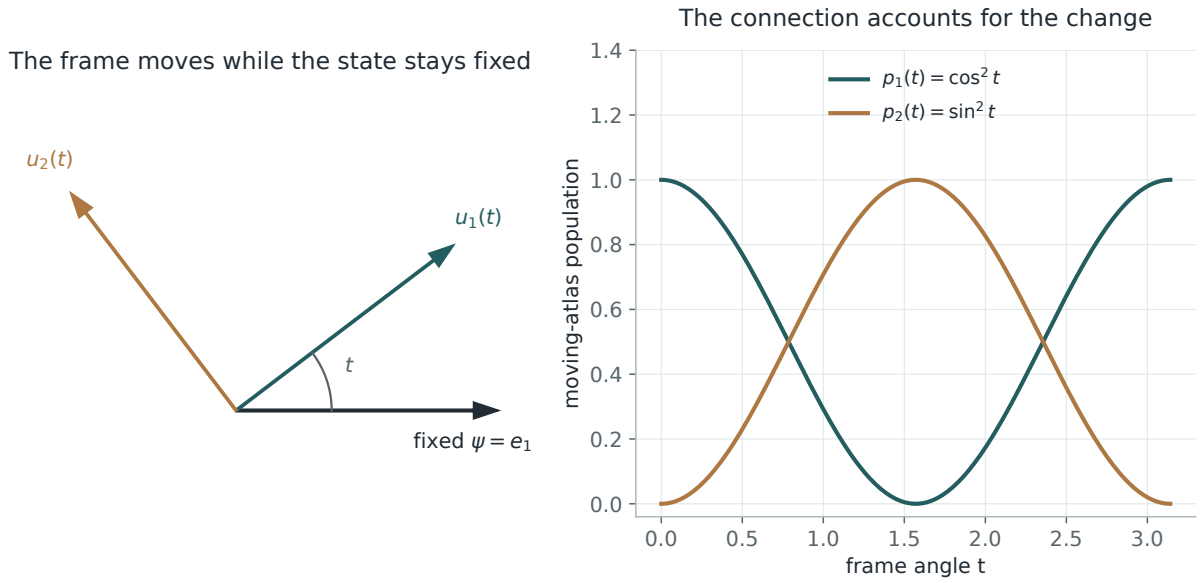


Figure 4. A stationary laboratory vector observed in a rotating atlas. The varying packet weights are generated by the connection term $-U^*\dot U$ in moving coordinates.

6 Mixed states and coherent transfer

For this section take $\mathcal{H} = \mathbb{C}^n$. A density matrix is a positive semidefinite Hermitian matrix ρ with $\operatorname{Tr}\rho = 1$. Its probability in a sharp packet is

$$p_a = \operatorname{Tr}(P_a\rho). \tag{19}$$

Hermiticity makes this trace real. Positivity makes it nonnegative, since cyclicity and $P_a^2 = P_a$ give $\operatorname{Tr}(P_a\rho) = \operatorname{Tr}(P_a\rho P_a) \geq 0$. Completeness gives $\sum_a p_a = \operatorname{Tr}\rho = 1$. These are separate uses of the density-matrix assumptions.

Suppose a differentiable matrix trajectory obeys the von Neumann equation

$$\dot\rho = -i[H,\rho], \qquad H = H^*. \tag{20}$$

For its coherent transfer define

$$J^{(\rho)}_{b\leftarrow a} = 2\operatorname{Im}\operatorname{Tr}(P_bHP_a\rho). \tag{21}$$

Proposition 6.1 (Trace-current balance)**.** *For a Hermitian ρ satisfying (20),*

$$\dot p_b = \sum_a J^{(\rho)}_{b\leftarrow a}, \qquad J^{(\rho)}_{b\leftarrow a} = -J^{(\rho)}_{a\leftarrow b}.$$

The subset boundary formula and exact coarse-current summation remain valid.

Proof. Let $z_{ba} = \operatorname{Tr}(P_bHP_a\rho)$. Taking the conjugate trace and cyclically moving the first factor to the end yields

$$\overline{z_{ba}} = \operatorname{Tr}(\rho P_aHP_b) = \operatorname{Tr}(P_aHP_b\rho) = z_{ab}.$$

Completeness and cyclicity resolve the derivative as

$$\dot p_b = -i\sum_a (z_{ba} - z_{ab}) = \sum_a 2\operatorname{Im} z_{ba}.$$

The conjugate relation gives antisymmetry. Finite sums then prove cancellation and coarse aggregation exactly as before. □

For a rank-one state $\rho = |\psi\rangle\langle\psi|$, where the associated operator sends x to $\psi\,\langle\psi, x\rangle$, one has $\operatorname{Tr}(M\rho) = \langle\psi, M\psi\rangle$. Substituting $M = P_bHP_a$ proves that (21) is precisely the pure-state formula. No normalization is needed for that identity itself.

In the coordinate atlas, the current is $2\operatorname{Im}(H_{ba}\rho_{ab})$. It depends on off-diagonal coherence. A useful test is to compare

$$\rho_+ = \frac{1}{2}\begin{pmatrix} 1 & i \\ -i & 1 \end{pmatrix}, \qquad \rho_- = \frac{1}{2}\begin{pmatrix} 1 & -i \\ i & 1 \end{pmatrix}, \qquad \rho_0 = \frac{1}{2}\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

All three are density matrices with packet probabilities $(1/2, 1/2)$. For $H = g\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$, however, their currents $J_{2\leftarrow 1}$ are respectively $g, -g, 0$. The same instantaneous populations can therefore have different derivatives under the same Hamiltonian. A closed equation involving only those two populations cannot represent all three coherent states.

The trace balance requires Hermiticity, but its algebra does not require positivity or trace one. Those additional hypotheses justify the probabilistic interpretation. This distinction is useful in formalization because it reveals which assumptions belong to a calculation and which belong to interpreting its output. Infinite-dimensional mixed states would require a trace-class setting and further control of products and derivatives. They are not silently included in this finite matrix argument.

7 Tensor packets, marginals, and retained correlations

Let $\mathcal{H}_A$ and $\mathcal{H}_B$ be finite-dimensional Hilbert spaces with finite orthogonal atlases (P_a) and (Q_μ). Their tensor atlas is

$$R_{a\mu} = P_a \otimes Q_\mu.$$

It is orthogonal because tensor multiplication gives $R_{a\mu}R_{b\nu} = \delta_{ab}\delta_{\mu\nu}R_{a\mu}$. Its sum is $\mathrm{id}_A \otimes \mathrm{id}_B$. For a density matrix ρ_{AB}, define

$$p_{a\mu} = \mathrm{Tr}(R_{a\mu}\rho_{AB}), \qquad p^A_a = \sum_\mu p_{a\mu}, \quad p^B_\mu = \sum_a p_{a\mu}. \tag{22}$$

These marginalizations are exact atlas coarsenings. For example,

$$p^A_a = \mathrm{Tr}((P_a \otimes \mathrm{id}_B)\rho_{AB}) = \mathrm{Tr}(P_a\rho_A), \qquad \rho_A = \mathrm{Tr}_B\,\rho_{AB}.$$

Choose full orthonormal bases of the two Hilbert spaces, with $(e_k)_k$ indexing the basis of $\mathcal{H}_B$. In the resulting product basis the partial trace has entries $(\rho_A)_{ij} = \sum_k (\rho_{AB})_{(i,k),(j,k)}$. This explicit formula shows that it preserves Hermiticity and trace. Positivity follows by evaluating the quadratic form on x and summing the nonnegative forms of $x \otimes e_k$. Thus the reduced matrix is itself a density matrix.

For a product state $\rho_{AB} = \sigma_A \otimes \tau_B$, the trace factorizes and $p_{a\mu} = p^A_a p^B_\mu$. General states need not have that property. The normalized vector $\Psi = (e_0 \otimes e_0 + e_1 \otimes e_1)/\sqrt{2}$ has joint probabilities

$$(p_{a\mu}) = \begin{pmatrix} 1/2 & 0 \\ 0 & 1/2 \end{pmatrix}, \qquad p^A = p^B = (1/2, 1/2).$$

The product density matrix $(\mathrm{id}_A/2) \otimes (\mathrm{id}_B/2)$ has the same marginals but every joint entry is $1/4$. Marginals therefore fail to reconstruct even the joint measurement distribution. Moreover, the classical mixture of $e_0 \otimes e_0$ and $e_1 \otimes e_1$ has the same joint table as Ψ but lacks its off-diagonal coherence. A single atlas table does not reconstruct the full density matrix or certify entanglement on its own.

When a joint coherent current is defined on the product labels, the current into a marginal packet is obtained by summing over the hidden labels on both its source and destination. Internal exchange within a marginal cell cancels. Correlations can still enter the surviving current, so exact marginal balance does not imply autonomous marginal dynamics. The distinction repeats the one already encountered in coarse pure-state transfer.

8 Open dynamics and rank-one jump flow

Coherent exchange is one mechanism for changing packet weights. A second comes from an explicitly chosen family of dissipative transitions. Fix an orthonormal basis $(e_a)_{a\in I}$ of $\mathbb{C}^n$, use its rank-one projectors, and choose rates $\gamma_{ba} \geq 0$ for $b \neq a$. Set diagonal rates to zero for convenience and define

$$L_{ba} = \sqrt{\gamma_{ba}}\,|e_b\rangle\langle e_a|, \qquad \mathcal{D}_L(\rho) = L\rho L^* - frac12(L^*L\rho + \rho L^*L). \tag{23}$$

Consider a supplied differentiable matrix trajectory satisfying

$$\dot\rho = \sum_{b\neq a} \mathcal{D}_{L_{ba}}(\rho). \tag{24}$$

This choice of jumps is tied to the selected basis. A general dissipative operator or a higher-rank atlas does not automatically produce the same closed rate equation.

Proposition 8.1 (Population gain and loss)**.** *The diagonal populations $p_b = \rho_{bb}$ of a Hermitian solution satisfy*

$$\dot p_b = \sum_a F_{b\leftarrow a}, \qquad F_{b\leftarrow a} = \gamma_{ba}p_a - \gamma_{ab}p_b, \qquad F_{b\leftarrow a} = -F_{a\leftarrow b}. \tag{25}$$

The total population is constant, and every coarse population has the sum of the boundary jump currents as its derivative.

Proof. For the one transition $a \to b$,

$$L_{ba}\rho L_{ba}^* = \gamma_{ba}p_aP_b, \qquad L_{ba}^*L_{ba} = \gamma_{ba}P_a.$$

Its diagonal gain is $\gamma_{ba}p_a$ at b, and the two halves of the anticommutator remove that same amount at a. No other diagonal entry changes from this one jump. Sum over ordered transitions to obtain (25). Antisymmetry is immediate, and the finite boundary cancellation applies. □

In this model the population equation closes without knowing the coherences. The coherences nevertheless have their own dynamics. For $i \neq j$, let $r_i = \sum_{b\neq i}\gamma_{bi}$ be the total rate leaving i. The gain matrices are diagonal, while each anticommutator selects its source row and column. Consequently

$$\dot\rho_{ij} = -\tfrac12(r_i + r_j)\rho_{ij}. \tag{26}$$

For fixed rates, $\rho_{ij}(t) = e^{-(r_i+r_j)t/2}\rho_{ij}(0)$. This decay law belongs to the stated rank-one jump family. Additional Hamiltonian or dephasing terms alter it in identifiable ways.

Example 8.2 (Two-state relaxation)**.** Take $\gamma_{21} = \gamma_{12} = k > 0$ and $p_1 + p_2 = 1$. Then

$$\dot p_1 = k(1-2p_1), \qquad p_1(t) = \tfrac12 + (p_1(0) - \tfrac12)e^{-2kt}, \qquad \rho_{12}(t) = e^{-kt}\rho_{12}(0).$$

The populations relax toward equal weights. In the coherent exchange of Example 3.2, the populations oscillate and depend on phase. Both systems admit antisymmetric probability currents, but that common bookkeeping does not make their mechanisms identical. Adding $-i[H,\rho]$ to (24) adds the coherent current to the jump current. The resulting population equation then generally depends on coherence again.

The coarse jump balance is always exact. For it to close as a time-homogeneous Markov equation for coarse weights alone, a further compatibility condition suffices. For distinct coarse cells A, B, suppose

$$\sum_{b \in C_B} \gamma_{ba} = q_{BA} \quad \text{for every } a \in C_A.$$

Then summing the fine gains gives $q_{BA} p^{\mathrm{c}}_A$, so $\dot{p}^{\mathrm{c}}_B = \sum_{A \neq B} (q_{BA} p^{\mathrm{c}}_A - q_{AB} p^{\mathrm{c}}_B)$. Without this equality of outgoing rates to each destination cell, the fine population distribution inside a source cell matters. This direct consequence of the gain sum shows precisely what an autonomous coarse model asks in addition to current conservation.

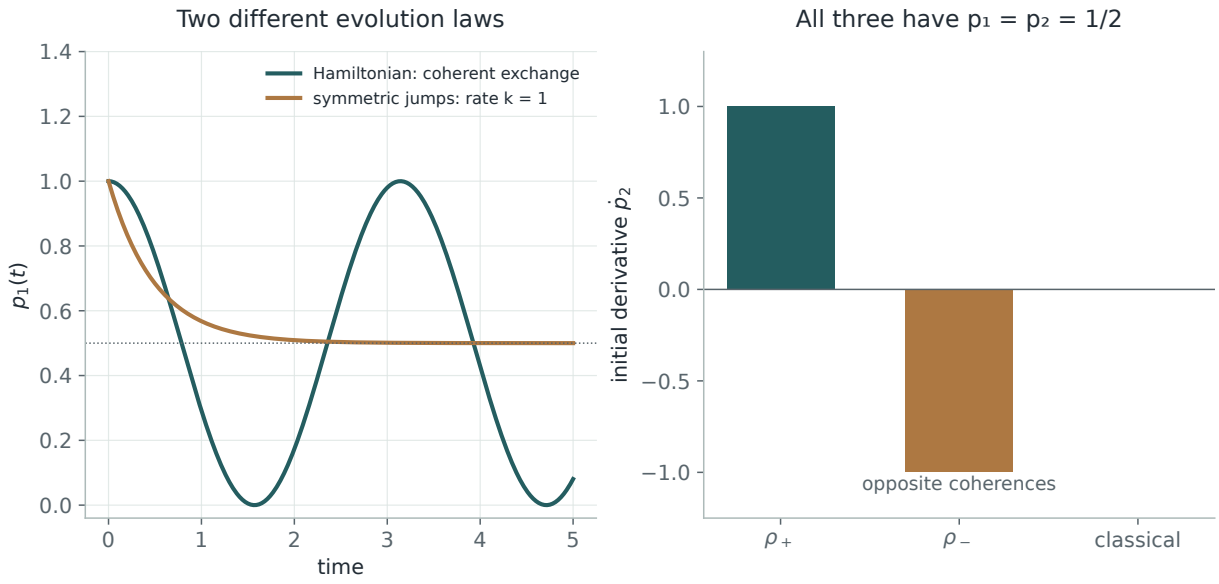


Figure 5. Coherent exchange and the specified symmetric rank-one jump dynamics have different population laws. The equal-population coherent states ρ_+ and ρ_- give opposite initial derivatives of p_2, while the classical jump equilibrium has zero derivative.

9 Information functionals and the role of detailed balance

An antisymmetric current also organizes derivatives of scalar functionals of the packet distribution. Let $p_a(t) > 0$ be differentiable, normalized, and obey $\dot{p}_b = \sum_a J_{b \leftarrow a}$ with J antisymmetric. For a differentiable real function f on the positive half-line,

$$\frac{d}{dt} \sum_b f(p_b) = \frac{1}{2} \sum_{b,a} J_{b \leftarrow a} (f'(p_b) - f'(p_a)). \tag{27}$$

To prove it, apply the ordinary chain rule and substitute the balance equation. Exchange the two indices in a second copy of the sum and use antisymmetry. Averaging the two expressions yields (27).

For Shannon entropy $S(p) = -\sum_a p_a \log p_a$, this gives

$$\dot{S} = -\frac{1}{2} \sum_{b,a} J_{b \leftarrow a} \log \frac{p_b}{p_a}. \tag{28}$$

The constant derivative term cancels because the total current is zero. For a real q, the power sum $M_q = \sum_a p_a^q$ has derivative

$$\dot{M}_q = \frac{q}{2} \sum_{b,a} J_{b \leftarrow a} (p_b^{q-1} - p_a^{q-1}).$$

For $q \neq 1$, the Rényi entropy $H_q = \log M_q/(1-q)$ therefore satisfies $\dot{H}_q = \dot{M}_q/((1-q)M_q)$. Strict positivity makes the powers and logarithmic derivatives valid, and $M_q > 0$ when the finite index set is nonempty. Boundary points with zero populations require their own limiting argument and are not covered by these derivative statements as written.

These formulas alone impose no sign on entropy change. The coherent two-state orbit can increase and then decrease its packet entropy as the state moves through a cycle. Packet entropy depends on the chosen atlas, and it is different from the von Neumann entropy of the density operator. In particular, a pure state can have positive packet entropy.

For the rank-one jump model, detailed balance supplies a definite monotonicity theorem. Let $\pi_a > 0$ be a fixed probability distribution, and assume

$$\gamma_{ba} \pi_a = \gamma_{ab} \pi_b. \tag{29}$$

Define relative densities $u_a = p_a/\pi_a$ and conductances $c_{ba} = \gamma_{ba}\pi_a$. Then $c_{ba} = c_{ab} \geq 0$ and $F_{b \leftarrow a} = c_{ba}(u_a - u_b)$.

Theorem 9.1 (Relative-entropy dissipation)**.** *For a strictly positive solution of the finite jump balance and a fixed positive reference satisfying* (29)*,*

$$D(p\|\pi) = \sum_a p_a \log \frac{p_a}{\pi_a}, \tag{30}$$

$$\frac{d}{dt} D(p\|\pi) = -\frac{1}{2} \sum_{b,a} c_{ba}(u_b - u_a)(\log u_b - \log u_a) \leq 0.$$

Proof. Differentiate with π fixed. The terms without logarithms sum to zero. Pairing opposite currents leaves $\frac{1}{2}\sum_{b,a} F_{b \leftarrow a}(\log u_b - \log u_a)$. Insert the conductance expression for F. The logarithm is increasing on $(0, \infty)$, so each product $(u_b - u_a)(\log u_b - \log u_a)$ is nonnegative. Every conductance is nonnegative, proving the inequality. □

The conductance form includes inactive edges without evaluating a logarithm of a zero flux. When $c_{ba} = 0$, its contribution is simply zero. When $c_{ba} > 0$, a summand vanishes exactly when $u_b = u_a$. Thus zero instantaneous dissipation means that the relative density is constant on each connected component of the positive-conductance graph. If that graph is connected, normalization forces the constant to be one and $p = \pi$. This identifies the equality case. The derivative calculation by itself does not provide a convergence rate to equilibrium.

The finite calculus now has a common structure across pure states, mixed states, moving coordinates, tensor marginals, and the specified jump model. Each setting begins with actual mathematical objects and a stated derivative. Reconstruction produces the balance, antisymmetry gives cancellation, and an exact grouping preserves the resulting current. Extending this chain to positive effects, countable families, or unbounded generators requires new hypotheses at the step where the original finite proof used orthogonality, finite summation, or bounded application.

10 Refinement cancellation and the cost of extrapolation

Coarsening a packet atlas combines exact components of one description. A different use of refinement compares approximations made at several scales. The distinction matters. An algebraic cancellation of leading error terms requires an expansion in a common vector space, and a useful numerical result also requires control of the remaining error and of the sample perturbations.

Let E be a real normed vector space and let $f(h) \in E$ be defined for $0 < |h| < \delta$. Fix a refinement ratio $b > 1$. If $f(h)$ represents operators on different Hilbert spaces, comparison maps must first place them in a common space where the following additions are meaningful. The formulas alone do not construct those maps.

The denominators below are positive. Each level is a finite linear combination of the original samples at successively refined scales.

Definition 10.1 (The actual refinement tower)**.** Set $T_0 f(h) = f(h)$ and, for $m \geq 0$, define

$$T_{m+1} f(h) = \frac{b^{m+1} T_m f(h/b) - T_m f(h)}{b^{m+1} - 1}.$$

Proposition 10.2 (Exact response of a monomial)**.** *For $r \in \mathbb{N}$ and $v \in E$,*

$$T_m(h^r v) = c_{m,r} h^r v, \qquad c_{m,r} = \prod_{k=1}^{m} \frac{b^{k-r} - 1}{b^k - 1}.$$

The empty product equals one. In particular, $c_{m,0} = 1$, and $c_{m,r} = 0$ for $1 \leq r \leq m$.

Proof. At level zero the assertion is the definition. If it holds at level m, the next recursion multiplies the coefficient by

$$\frac{b^{m+1} b^{-r} - 1}{b^{m+1} - 1}.$$

This appends the next factor to the product. For $r = 0$, each factor equals one. For $1 \leq r \leq m$, the factor at $k = r$ vanishes. □

Suppose that a finite expansion is available,

$$f(h) = F + \sum_{r=1}^{m} h^r v_r + R(h).$$

Linearity and the proposition give the exact identity

$$T_m f(h) - F = T_m R(h).$$

Only a finite sum has been used. No convergence of an infinite formal series is needed. The size of the right-hand side remains an analytic question.

Theorem 10.3 (A quantified remainder and perturbation bound)**.** *Let p be a nonnegative integer and suppose*

$$\|R(h)\| \le C|h|^p, \qquad \|g(h) - f(h)\| \le \varepsilon \quad (0 < |h| < \delta),$$

where $C, \varepsilon \ge 0$. Define

$$A_{m,p} = \prod_{k=1}^{m} \frac{b^{k-p} + 1}{b^k - 1}.$$

Then, throughout the same punctured interval,

$$\|T_m g(h) - F\| \le A_{m,p} C|h|^p + A_{m,0}\varepsilon.$$

Proof. The interval is stable under $h \mapsto h/b$. The triangle inequality applied to one tower step gives

$$\|T_{m+1}R(h)\| \le \frac{b^{m+1}\,\|T_m R(h/b)\| + \|T_m R(h)\|}{b^{m+1} - 1}.$$

Starting with the stated remainder bound and proceeding by induction multiplies its constant at step $m+1$ by $(b^{m+1-p}+1)/(b^{m+1}-1)$. Hence $\|T_m R(h)\| \le A_{m,p}C|h|^p$. Apply the same argument with $p = 0$ to $g - f$. Linearity and the triangle inequality separate the two contributions. □

When $p = m+1$, this proves the expected order improvement under an explicit remainder hypothesis. It also retains the effect of noisy samples. Each factor in $A_{m,0}$ exceeds one, so extrapolation can amplify a uniform sample error even while cancelling several deterministic powers of h.

10.1 A complete three-level example

For $b = 2$, direct substitution gives

$$\begin{aligned} T_1 f(h) &= 2f(h/2) - f(h), \\ T_2 f(h) &= \frac{8f(h/4) - 6f(h/2) + f(h)}{3}, \\ T_3 f(h) &= \frac{64f(h/8) - 56f(h/4) + 14f(h/2) - f(h)}{21}. \end{aligned}$$

Take the exact scalar polynomial $f(h) = 1 + 2h - 3h^2 + 4h^3$. The responses are

$$T_1 f(h) = 1 + \frac{3}{2}h^2 - 3h^3, \qquad T_2 f(h) = 1 + \frac{1}{2}h^3, \qquad T_3 f(h) = 1.$$

These are polynomial identities for every real h. If each of the three samples used by T_2 has absolute error at most ε, its output error is at most

$$\frac{8 + 6 + 1}{3}\varepsilon = 5\varepsilon.$$

This bound is attained by errors with signs matching the three weights. Thus exact removal of the first two powers does not imply a contraction on arbitrary perturbations. For the displayed polynomial, the perturbed second-level approximation has the explicit total bound $|h|^3/2 + 5\varepsilon$.

The calculation supplies a practical reading rule for later refinement claims. An exact coefficient identity, an asymptotic remainder estimate, and a stability bound are separate pieces of evidence. A

physical approximation uses all three, together with the comparison maps that make its values commensurable.

Limits, domains, and differential structure

11 Positive effects and sharp dilation

A sharp packet atlas assigns mutually orthogonal subspaces to its labels. Many useful descriptions overlap. A detector may respond partly to several channels, and a frame may analyze the same component more than once. The correct replacement is a positive partition of the identity. Its weights are expectation values, and its channel vectors live in a larger space where orthogonality can be restored.

Throughout this chapter, inner products are conjugate linear in the first variable. A bounded state generator is denoted by K, with $K^* = -K$. When a Hamiltonian is present, $K = -iH$. The receiving label is written first in J_{ba}.

Definition 11.1 (Finite effect atlas)**.** Let $\mathcal{H}$ be a complex Hilbert space and let A be a finite label set. An effect atlas is a family of bounded operators satisfying

$$F_a = F_a^* \geq 0, \qquad \sum_{a \in A} F_a = I.$$

Its weight in a state ψ is $p_a = \langle \psi, F_a \psi \rangle$. The partition identity implies $0 \leq F_a \leq I$, so $p_a \geq 0$ and $\sum_a p_a = \|\psi\|^2$.

For a sharp atlas, $F_a = P_a$ and $p_a = \|P_a \psi\|^2$. In general, p_a is not $\|F_a \psi\|^2$. The latter quantity is $\langle \psi, F_a^2 \psi \rangle$ and therefore loses precisely the positive defect $F_a - F_a^2$.

Example 11.2 (Overlapping two-channel weights)**.** On $\mathbb{C}^2$, take

$$F_+ = \begin{pmatrix} 4/5 & 0 \\ 0 & 1/5 \end{pmatrix}, \qquad F_- = \begin{pmatrix} 1/5 & 0 \\ 0 & 4/5 \end{pmatrix}, \qquad \psi = \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

The true weights are $4/5$ and $1/5$. Squared effect-vector norms would instead give $16/25$ and $1/25$, whose sum is $17/25$. The state still satisfies $F_+ \psi + F_- \psi = \psi$, but its two summands are not orthogonal. Reconstruction and an orthogonal decomposition are different assertions.

11.1 Continuity without orthogonality

Suppose that a differentiable state curve satisfies $\dot{\psi} = K\psi$. Define

$$J_{ba}^F = \operatorname{Re} \langle F_b \psi, K F_a \psi \rangle - \operatorname{Re} \langle F_a \psi, K F_b \psi \rangle . \tag{31}$$

This current is antisymmetric directly from its definition. Skew-adjointness also gives

$$J_{ba}^F = 2 \operatorname{Re} \langle F_b \psi, K F_a \psi \rangle .$$

The same bilinear expression occurs for sharp packets, but its interpretation now concerns expectation weights.

Proposition 11.3 (Effect continuity)**.** *Under these hypotheses,*

$$\dot{p}_b = \sum_a J^F_{ba}, \qquad \sum_b \dot{p}_b = 0.$$

If $F_b = F_b(t)$ is differentiable in operator norm, the first identity becomes

$$\dot{p}_b = s_b + \sum_a J^F_{ba}, \qquad s_b = \left\langle \psi, \dot{F}_b \psi \right\rangle, \qquad \sum_b s_b = 0.$$

Proof. Self-adjointness of F_b and differentiation give

$$\dot{p}_b = 2\,\mathrm{Re}\,\langle F_b\psi, K\psi\rangle$$

for a fixed effect. Summing the two terms of (31) over a and using $\sum_a F_a = I$ gives

$$\sum_a J^F_{ba} = \mathrm{Re}\,\langle F_b\psi, K\psi\rangle - \mathrm{Re}\,\langle \psi, K F_b\psi\rangle\,.$$

The second real part is the negative of the first. The moving case adds the product-rule term s_b. Its total is zero because $\sum_b \dot{F}_b = 0$. Antisymmetry cancels the finite double sum of currents. □

If a map $r : A \to B$ groups the labels, the coarse effect is $G_\beta = \sum_{r(a)=\beta} F_a$. Bilinearity proves

$$p^{\mathrm{c}}_\beta = \sum_{r(a)=\beta} p_a, \qquad J^{\mathrm{c}}_{\beta\alpha} = \sum_{\substack{r(b)=\beta \\ r(a)=\alpha}} J^F_{ba}.$$

No orthogonality is needed for this finite aggregation law. A divergence-free antisymmetric array may also be added to J^F without changing its row sums. Thus the balance law determines the change in weight, while the generator-induced formula selects a particular representative of the current.

11.2 Analysis, synthesis, and the range projection

Choose Hilbert spaces $\mathcal{K}_a$ and bounded maps $V_a : \mathcal{H} \to \mathcal{K}_a$ such that

$$\sum_a V_a^* V_a = I.$$

The combined analysis map $V : \mathcal{H} \to \mathcal{K} = \bigoplus_a \mathcal{K}_a$ is $V\psi = (V_a\psi)_a$. Its adjoint is synthesis, $V^*(y_a)_a = \sum_a V_a^* y_a$. Consequently

$$V^*V = I, \qquad R = VV^*, \qquad R^* = R^2 = R.$$

Here R projects onto the closed range of the isometry V. The range need not fill $\mathcal{K}$. If Π_a is the coordinate projector, the compressed effect is

$$F_a = V^*\Pi_a V = V_a^* V_a, \qquad p_a = \|V_a\psi\|^2 = \|\Pi_a V\psi\|^2\,.$$

Every finite effect atlas admits this construction by taking $\mathcal{K}_a = \mathcal{H}$ and $V_a = F_a^{1/2}$. Positivity supplies the bounded positive square root. No choice of a measurement outcome or post-measurement state is being made by this algebraic dilation.

For $X \in \mathcal{B}(\mathcal{H})$, its channel block is $X^V_{ba} = V_b X V_a^*$. Inserting the partition on both sides proves

$$X = \sum_{a,b} V_b^* X^V_{ba} V_a. \tag{32}$$

The lifted operator $\widetilde{X} = VXV^*$ is multiplicative on lifted operators, since $V^*V = I$. Its identity is R, not $I_{\mathcal{K}}$. Conversely, compressing arbitrary ambient operators need not preserve multiplication:

$$V^*ABV - (V^*AV)(V^*BV) = V^*A(I-R)BV. \tag{33}$$

This is the exact missing-range term. Setting $A = B = \Pi_a$ gives

$$\langle \psi, (F_a - F_a^2)\psi \rangle = \left\| (I-R)\Pi_a V\psi \right\|^2. \tag{34}$$

Thus the failure of an effect to be sharp measures how much its coordinate projection leaves the embedded physical range. This identity also proves $F_a - F_a^2 \geq 0$ directly.

The dynamics respects the same construction. Put $\widetilde{K} = VKV^*$. It is skew-adjoint, and $\frac{d}{dt} V\psi = \widetilde{K} V\psi$. Moving adjoints across the inner product shows

$$2\,\mathrm{Re}\left\langle \Pi_b V\psi, \widetilde{K}\Pi_a V\psi \right\rangle = 2\,\mathrm{Re}\,\langle F_b\psi, KF_a\psi\rangle = J^F_{ba}.$$

The positive-effect current is therefore exactly the sharp current of this lifted flow evaluated on the embedded state. It is not obtained by treating arbitrary compression as an algebra homomorphism.

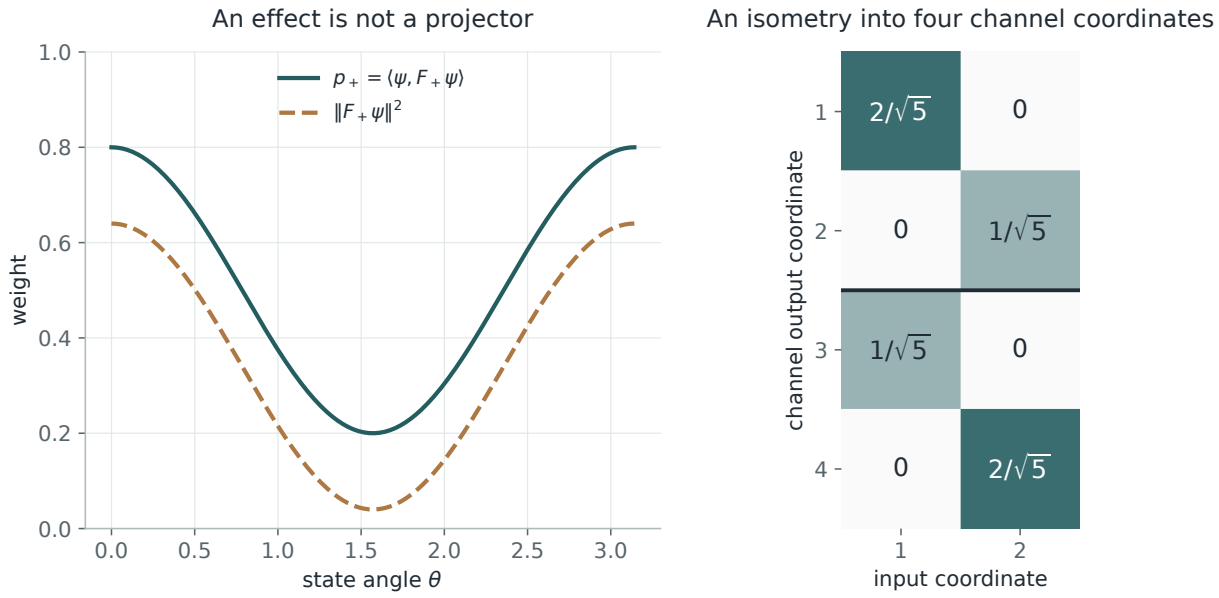


Figure 6. An explicit two-effect example with $F_+ = \mathrm{diag}(4/5, 1/5)$ and $F_- = I - F_+$. The square-root channel isometry preserves the effect probabilities in an enlarged sharp resolution.

12 Countable packets and the domain of a generator

An infinite resolution introduces two separate limits. Packet vectors must reconstruct the state, and applying an unbounded operator to the partial reconstruction requires additional control. Writing an infinite sum of matrix blocks conceals this distinction unless the domain is stated first.

Let $(P_a)_{a\in A}$ be a countable orthogonal family of bounded projections with strong sum I. For finite $F \subset A$, write $S_F = \sum_{a\in F} P_a$. Strong completeness means $S_F x \to x$ for each fixed $x \in \mathcal{H}$, over the directed set of finite subsets. Orthogonality gives

$$\|S_F x\|^2 = \sum_{a\in F} \|P_a x\|^2, \qquad \sum_a \|P_a x\|^2 = \|x\|^2.$$

These are Hilbert-space statements. They contain no unbounded operator.

Now let $H : D(H) \subset \mathcal{H} \to \mathcal{H}$ be a densely defined symmetric operator. To write $HP_a\psi$ for every domain state, require

$$P_a D(H) \subset D(H). \tag{35}$$

Define $\|x\|_H^2 = \|x\|^2 + \|Hx\|^2$. A domain-compatible atlas is graph compatible when

$$S_F x \longrightarrow x \quad \text{in the graph norm for every } x \in D(H). \tag{36}$$

Equivalently, both $S_F x \to x$ and $HS_F x \to Hx$ hold in $\mathcal{H}$. The second convergence does not follow from the first for an unbounded H.

Example 12.1 (A small state can have a large generator image)**.** On $\ell^2(\mathbb{N}_{\geq 1})$, let $He_n = ne_n$ on its maximal domain. The finitely supported vectors $x_N = e_N/N$ satisfy $\|x_N\| = 1/N \to 0$, while $\|Hx_N\| = 1$. This gives a direct reason to track the generator tail in addition to the state tail. It does not contradict closedness, because $Hx_N = e_N$ has no norm limit.

12.1 A maximal diagonal domain

Let (λ_n) be any real sequence and define

$$D(H_\lambda) = \left\{ c \in \ell^2 : \sum_n \lambda_n^2 |c_n|^2 < \infty \right\}, \qquad (H_\lambda c)_n = \lambda_n c_n.$$

The operator is dense because it contains every finitely supported sequence. It is symmetric by the reality of the multipliers. Its domain is maximal for this coordinate action.

Proposition 12.2 (Self-adjointness of the diagonal realization)**.** *The operator H_λ is self-adjoint. Its coordinate cutoffs converge in graph norm on $D(H_\lambda)$.*

Proof. If $y \in D(H_\lambda^*)$ and $H_\lambda^* y = z$, testing the adjoint identity against e_n gives $z_n = \lambda_n y_n$. Since $z \in \ell^2$, the sequence y belongs to $D(H_\lambda)$. Symmetry gives the opposite inclusion of domains, and the actions agree there. For $S_N c = \sum_{n\leq N} c_n e_n$,

$$\|c - S_N c\|_{H_\lambda}^2 = \sum_{n>N} (1 + \lambda_n^2)|c_n|^2 \longrightarrow 0.$$

□

For positive λ_n, the quadratic-form domain instead uses $\sum_n \lambda_n |c_n|^2 < \infty$. This is the domain of $H_\lambda^{1/2}$, not generally the domain of H_λ. With $\lambda_n = n$ and $c_n = n^{-s}$, the three thresholds are

membership	condition
ℓ^2	$s > 1/2$
$D(H_\lambda^{1/2})$	$s > 1$
$D(H_\lambda)$	$s > 3/2.$

Each equality case diverges harmonically. Thus a state with finite mean positive energy can still lack the Hilbert-space derivative required by the Schrödinger equation. The harmonic oscillator example has $\lambda_n = n + 1/2$ for $n \geq 0$ and the same asymptotic thresholds after replacing n by $n + 1$ in the coefficients.

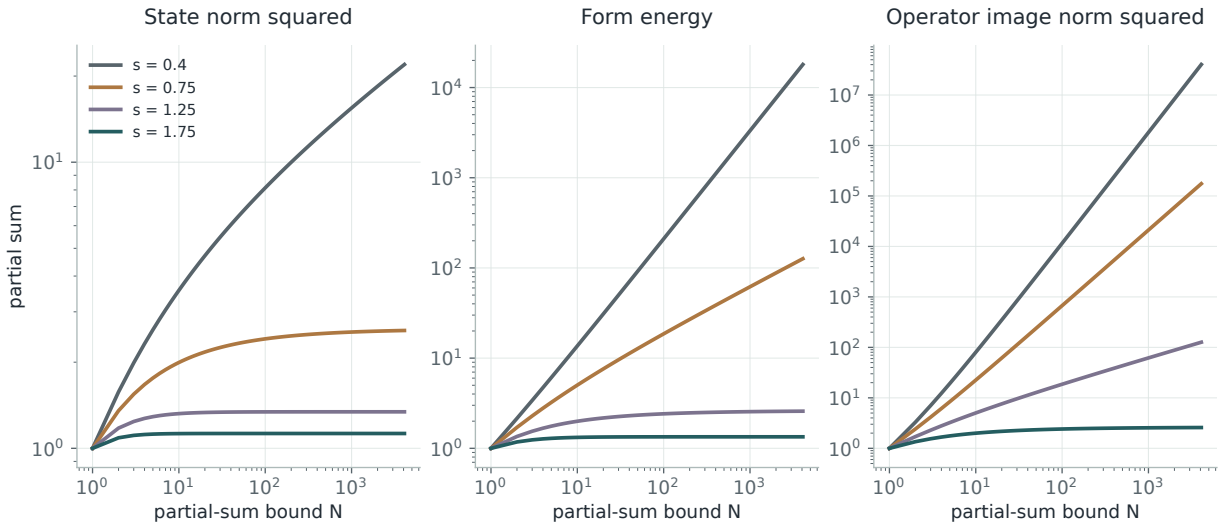


Figure 7. For $He_n = ne_n$, $n \geq 1$, coefficients $a_n = n^{-s}$ encounter three different thresholds: state norm, quadratic-form domain, and operator domain. The finite partial-sum curves illustrate the analytic p-series criteria.

12.2 Countable continuity and the stronger scalar-row argument

Suppose a supplied curve $\psi(t) \in D(H)$ is differentiable as a curve in $\mathcal{H}$ and satisfies $\dot{\psi} = -iH\psi$. For a domain-preserving atlas, put

$$p_b = \|P_b\psi\|^2, \qquad J_{ba} = 2\,\mathrm{Im}\,\langle P_b\psi, HP_a\psi\rangle\,.$$

Symmetry gives $J_{ba} = -J_{ab}$. Differentiating the bounded projected norm gives

$$\dot{p}_b = 2\,\mathrm{Im}\,\langle P_b\psi, H\psi\rangle\,.$$

Graph compatibility permits direct insertion of $HS_F\psi \to H\psi$. The formal development also proves a stronger statement for one receiving packet: symmetry moves H onto that packet, so only the strong state sum is needed. Indeed,

$$\langle P_b\psi, HP_a\psi\rangle = \langle HP_b\psi, P_a\psi\rangle\,,$$

and the sum over a converges to $\langle HP_b\psi, \psi\rangle = \langle P_b\psi, H\psi\rangle$. In fact the row is absolutely summable, since

$$\sum_a |\langle HP_b\psi, P_a\psi\rangle| = \sum_a |\langle P_aHP_b\psi, P_a\psi\rangle|$$
$$\leq \left(\sum_a \|P_aHP_b\psi\|^2\right)^{1/2} \left(\sum_a \|P_a\psi\|^2\right)^{1/2} = \|HP_b\psi\|\,\|\psi\|\,.$$

Consequently $\dot{p}_b = \sum_a J_{ba}$ without assuming graph compatibility. The two tail estimates describe the distinct arguments:

$$\left|\sum_{a\in F} J_{ba} - \dot{p}_b\right| \leq 2\,\|P_b\psi\|\,\|HS_F\psi - H\psi\|\,, \tag{37}$$

$$\left|\sum_{a\in F} J_{ba} - \dot{p}_b\right| \leq 2\,\|HP_b\psi\|\,\|S_F\psi - \psi\|\,. \tag{38}$$

The first bound uses graph-action control. The second uses symmetry and a fixed receiving packet. Neither states that the entire double array (J_{ba}) is absolutely summable.

12.3 Infinite coarse cells

For a partition $A = \bigsqcup_\alpha C_\alpha$, the strong sum $Q_\alpha = \sum_{a\in C_\alpha} P_a$ is a bounded orthogonal projection. Its weight is always

$$\|Q_\alpha\psi\|^2 = \sum_{a\in C_\alpha} \|P_a\psi\|^2\,.$$

To define a coarse current with $HQ_\alpha\psi$, additionally require $Q_\alpha D(H) \subset D(H)$. Fine domain preservation alone does not state this infinite-cell property.

Under this coarse-domain hypothesis, symmetry and the strong cell sums justify both iterated sums separately:

$$J^{\mathrm{c}}_{\beta\alpha} = 2\,\mathrm{Im}\,\langle Q_\beta\psi, HQ_\alpha\psi\rangle$$
$$= \sum_{b\in C_\beta}\left(\sum_{a\in C_\alpha} J_{ba}\right) = \sum_{a\in C_\alpha}\left(\sum_{b\in C_\beta} J_{ba}\right).$$

For example, first summing a moves H onto $P_b\psi$, then the outer sum pairs $\sum_b P_b\psi$ with the fixed vector $HQ_\alpha\psi$. The reverse order has its own symmetric argument. Flattening these expressions to a single unordered sum on $C_\beta \times C_\alpha$ requires joint summability. Equality of two justified iterated sums does not supply that hypothesis.

For a simple warning, the antisymmetric scalar array with $a_{2n,2n+1} = 1/n$, $a_{2n+1,2n} = -1/n$ for $n \geq 1$, and all other entries zero has finite support in each row but infinite total absolute sum. This array is only a summation example, not an asserted physical realization of the current. It shows why rowwise convergence alone cannot authorize arbitrary rearrangement.

The coarse balance is obtained by differentiating $\|Q_\beta\psi(t)\|^2$, using the bounded coarse projector. There is no need to interchange an infinite sum of derivatives. Likewise, the total packet norm is identified with $\|\psi(t)\|^2$ by Parseval before its conservation is established. The order of these arguments keeps the domain and summation obligations visible.

13 Constructing evolution and controlling truncation

A continuity theorem that assumes a differentiable solution does not construct one. For the real diagonal model, the construction is explicit and the precise differentiability domain can be proved. This supplies a complete small model in which the preceding distinctions can be inspected directly.

For real (λ_n) define

$$(U(t)c)_n = e^{-it\lambda_n}c_n, \qquad c \in \ell^2. \tag{39}$$

The modulus of every multiplier is one. Hence $U(t)$ is a unitary operator, $U(t+s) = U(t)U(s)$, and $U(-t) = U(t)^*$. For each fixed state it is strongly continuous. To prove the latter claim, split c into a finite head and a small ℓ^2 tail. The finite head is continuous coordinatewise, while the norm of the difference on the tail is at most twice its initial norm.

Proposition 13.1 (The exact differentiability domain)**.** *The orbit $t \mapsto U(t)c$ is differentiable in ℓ^2 at one, equivalently every, time precisely when $c \in D(H_\lambda)$. On that domain,*

$$\frac{d}{dt}U(t)c = -iH_\lambda U(t)c, \qquad H_\lambda U(t)c = U(t)H_\lambda c.$$

Proof. If $c \in D(H_\lambda)$, each coordinate difference quotient tends to $-i\lambda_n e^{-it\lambda_n}c_n$. The elementary bound $|e^{-iu} - 1| \le |u|$ bounds the quotient magnitude by $|\lambda_n c_n|$. The squared difference from its limit is bounded by $4\lambda_n^2|c_n|^2$, a summable sequence. Dominated convergence therefore gives convergence in ℓ^2.

Conversely, if an ℓ^2 derivative d exists, bounded coordinate evaluation gives

$$d_n = -i\lambda_n e^{-it\lambda_n}c_n.$$

Since $d \in \ell^2$, one has $\sum_n \lambda_n^2|c_n|^2 < \infty$. Commutation and domain preservation follow from the same coordinate formula. □

The constructed orbit is the unique global domain-valued solution with a given domain initial state. If two supplied solutions have the same initial value, their difference w satisfies $\dot{w} = -iH_\lambda w$. Symmetry gives

$$\frac{d}{dt}\|w\|^2 = 2\,\mathrm{Re}\,\langle w, -iH_\lambda w\rangle = 0.$$

The initial zero norm remains zero. This uniqueness proof extends to supplied solutions of a symmetric operator, but the proof of existence above uses the explicit real diagonal construction.

For domain states, the orbit also satisfies the quantitative bound

$$\|U(t)c - U(s)c\| \le |t-s|\,\|H_\lambda c\|.$$

The constant belongs to the state. It is not an operator-norm bound for the unbounded generator. Spectral coordinate populations remain constant and their interpacket currents vanish. A different domain-preserving atlas can mix spectral coordinates and carry nonzero currents under the same evolution.

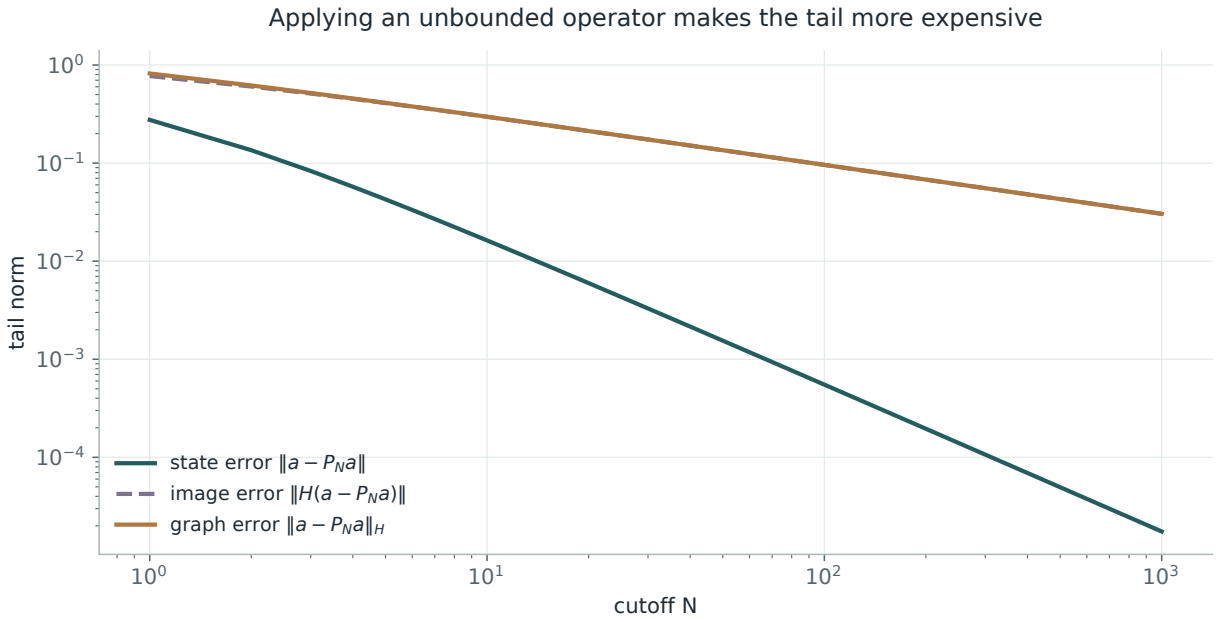


Figure 8. For the normalized state proportional to n^{-2} and $He_n = ne_n$, the state tail and generator tail decay at different rates. Graph control retains both contributions.

13.1 Uniformity in time for a fixed state

Coordinate cutoffs commute with $U(t)$. Thus for every $c \in \ell^2$,

$$\|U(t)c - S_N U(t)c\| = \|c - S_N c\|\,. \tag{40}$$

The right side tends to zero independently of $t \in \mathbb{R}$. This gives uniform convergence over all times for each fixed initial state. For $c \in D(H_\lambda)$, differentiating the finite cutoff orbit and the full orbit gives a second exact identity:

$$\left\| \frac{d}{dt} U(t)c - \frac{d}{dt} S_N U(t)c \right\| = \|H_\lambda c - S_N H_\lambda c\|\,. \tag{41}$$

This tail also tends to zero uniformly in time. State convergence alone does not imply this velocity convergence.

Example 13.2 (Strong convergence with no differentiable limiting orbit)**.** For $\lambda_n = n$, take $c_n = C/n$ with $C = (\sum_{n\geq 1} n^{-2})^{-1/2}$. This is a normalized ℓ^2 state outside $D(H_\lambda)$. Every finite cutoff orbit is differentiable, and (40) gives uniform convergence to the full orbit. Yet the full orbit is nowhere differentiable in ℓ^2, by the preceding proposition. In fact the norm of the cutoff velocity is $C\sqrt{N}$, which diverges.

For the faster normalized tail $c_n = C/n^2$, both errors are controlled:

$$\|c - S_N c\|^2 = C^2 \sum_{n>N} n^{-4}, \qquad \|H_\lambda c - S_N H_\lambda c\|^2 = C^2 \sum_{n>N} n^{-2}.$$

The state tail is $O(N^{-3/2})$ in norm, while the velocity tail is $O(N^{-1/2})$. A small state error can therefore coexist with a much larger derivative error even inside the true domain.

14 Which topology does reconstruction use?

Let S_N be orthogonal projections tending strongly to I. Their ranks need not be finite. For a bounded operator X, the reconstructed corner sum is $S_N X S_N$. Its natural convergence can be stronger when X belongs to a smaller operator class.

Theorem 14.1 (Three reconstruction statements)**.** *Under these hypotheses:*

1. *If X is bounded, $S_N X S_N \to X$ strongly.*
2. *If X is compact, $\|S_N X S_N - X\| \to 0$ in operator norm.*
3. *If X is Hilbert–Schmidt, $\|S_N X S_N - X\|_2 \to 0$ in Hilbert–Schmidt norm.*

Proof. Use the identity

$$S_N X S_N - X = S_N X (S_N - I) + (S_N - I) X.$$

Since $\|S_N\| \le 1$, application to a fixed x bounds the error by

$$\|X\| \, \|(S_N - I)x\| + \|(S_N - I)Xx\| \,,$$

which tends to zero.

For compact X, choose finite-rank Y with $\|X - Y\| < \varepsilon$. Strong convergence is uniform on each fixed finite-dimensional subspace, and self-adjointness of S_N controls the right error by the corresponding left error for Y^*. Thus $\|S_N Y S_N - Y\| \to 0$. The difference between the two reconstruction errors is at most 2ε, proving the claim.

For the last assertion, fix an orthonormal basis (e_j) and define

$$\|X\|_2^2 = \sum_j \|Xe_j\|^2 < \infty.$$

Each column $(S_N - I)Xe_j$ tends to zero, and its squared norm is bounded by $4\|Xe_j\|^2$. Dominated convergence gives $\|(S_N - I)X\|_2 \to 0$. The right error follows by taking adjoints, using $S_N^* = S_N$ and $\|X^*\|_2 = \|X\|_2$. The ideal inequality $\|AY\|_2 \le \|A\| \, \|Y\|_2$ and the displayed decomposition finish the proof. □

The Hilbert–Schmidt norm is independent of the chosen orthonormal basis, and it bounds the operator norm. It measures the total square size of all matrix columns. In a coordinate packet basis, it also measures the square sum of all matrix entries. Retaining a larger square of entries removes a correspondingly smaller summable tail. Operator norm instead measures the largest action on a unit vector, allowing the testing vector to depend on the cutoff.

Example 14.2 (The identity never has a small omitted tail in operator norm)**.** On an infinite-dimensional coordinate space, let S_N retain the first N coordinates. For $X = I$,

$$\|S_N I S_N - I\| = \|S_N - I\| = 1.$$

A unit vector e_{N+1} witnesses the lower bound. Nevertheless, $S_N x \to x$ for every fixed $x \in \ell^2$. This is the difference between pointwise control over states and uniform control over the full unit ball.

14.1 Zero padding and the correct asymptotic tail

For a diagonal bounded multiplier $f(H_\lambda)$, zero padding gives the exact formula

$$\|f(H_\lambda) - S_N f(H_\lambda) S_N\| = \sup_{n>N} |f(\lambda_n)|.$$

For the unitary multiplier $f(\lambda) = e^{-it\lambda}$, this norm is one at every finite cutoff. Equation (40) remains valid. Uniformity in time for a fixed state and failure of operator-norm convergence are therefore compatible.

For $\lambda_n = n$, a resolvent at $i\eta$, $\eta \neq 0$, has zero-padded error

$$\sup_{n>N} \frac{1}{\sqrt{n^2+\eta^2}} = \frac{1}{\sqrt{(N+1)^2+\eta^2}} \longrightarrow 0.$$

The heat multiplier e^{-tn} has error $e^{-t(N+1)}$ for fixed $t > 0$. These tails decay because the operators are compact. At $t = 0$ the heat operator is the identity, so the same norm-convergence assertion is false.

Two other bounded multipliers approach one along this positive spectrum:

$$b(n) = \frac{n}{\sqrt{1+n^2}}, \qquad q(n) = \frac{n-i}{n+i}.$$

Their correct anchored cutoff is

$$A_N(f) = S_N f(H_\lambda) S_N + (I - S_N).$$

The errors are then

$$\|b(H_\lambda) - A_N(b)\| = \sup_{n>N} \left| \frac{n}{\sqrt{1+n^2}} - 1 \right| \longrightarrow 0,$$

and

$$\|q(H_\lambda) - A_N(q)\| = \sup_{n>N} \frac{2}{\sqrt{n^2+1}} \longrightarrow 0.$$

The identity tail is an analytic choice justified by these asymptotic multipliers. It is not a universal prescription for every two-sided or non-diagonal spectrum.

On a graph-compatible atlas, an unbounded generator has another reconstruction statement:

$$S_N H S_N \psi - H\psi = S_N(HS_N\psi - H\psi) + (S_N - I)H\psi \longrightarrow 0, \qquad \psi \in D(H).$$

This is convergence of the displayed output vectors in the ambient Hilbert norm. It neither takes an operator norm of H nor asserts graph-norm convergence of those output vectors, which would require additional domain and second-action information.

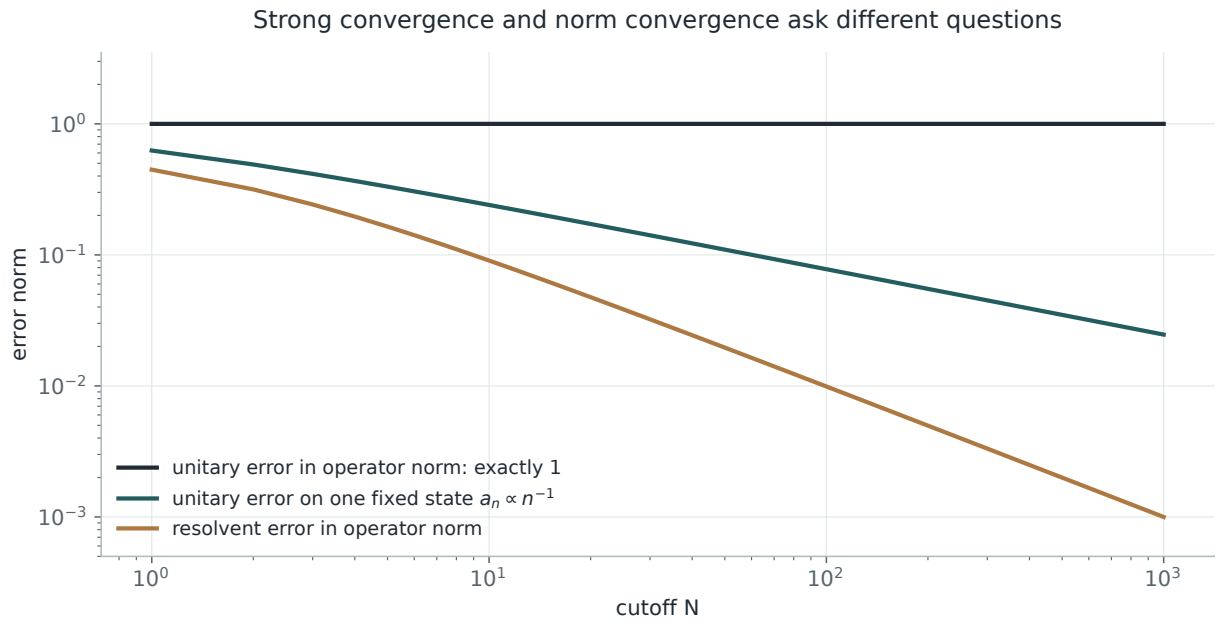


Figure 9. Zero-padded unitary cutoffs converge on a fixed state but retain operator-norm error one. For the same spectrum $n \geq 1$, the resolvent multiplier has a tail that decays in operator norm.

14.2 A bounded certificate for a nonreducing packet

A packet need not commute with an unbounded Hamiltonian. Suppose now that H is self-adjoint, P is a bounded orthogonal projection preserving $D(H)$, and $[H, P]$ on $D(H)$ extends to a bounded operator B. For nonreal z, let $R_z = (H - z)^{-1}$. It maps $\mathcal{H}$ into $D(H)$, so all products used below have their stated meaning. Multiplying the commutator identity on the domain by the resolvents gives

$$[R_z, P] = -R_z B R_z.$$

Indeed, for $x \in \mathcal{H}$, apply $H - z$ to $PR_z x$ to obtain $Px + BR_z x$, then apply R_z. Hence

$$\|[R_z, P]\| \leq \|R_z\|^2 \|B\|, \qquad \|[R_{i\eta}, P]\| \leq \frac{\|B\|}{\eta^2} \quad (\eta \neq 0).$$

The estimate converts a domain-sensitive commutator into a bounded one. In the generic formal certificate, the inverse and its domain properties are supplied explicitly. The diagonal modules also construct concrete inverses. Neither statement licenses a formal inverse for an unspecified operator.

15 The finite differential obstruction and matrix corners

Finite arrays support many useful difference operators. What they cannot support is a nonzero linear derivation of the full pointwise algebra $\mathbb{C}^N$ with its ordinary product. This obstruction explains why a finite operator calculus must declare which algebra carries its differentiation rule.

Theorem 15.1 (No nonzero derivation of the finite pointwise algebra)**.** *Let $D : \mathbb{C}^N \to \mathbb{C}^N$ be complex linear and satisfy*

$$D(fg) = D(f)g + fD(g)$$

for every pair of vectors, where multiplication is coordinatewise. Then $D = 0$.

Proof. For the coordinate idempotent e_a, the identity $e_a^2 = e_a$ gives $D(e_a) = 2e_aD(e_a)$. At a coordinate $b \neq a$, the right side is zero. At $b = a$, the equality reads $(D(e_a))_a = 2(D(e_a))_a$, again forcing zero. Thus $D(e_a) = 0$ for every basis vector, and linearity gives $D = 0$. □

The theorem concerns the semisimple pointwise algebra $\mathbb{C}^N$. It does not assert that every finite-dimensional commutative complex algebra has only the zero derivation. For example, on $\mathbb{C}[\varepsilon]/(\varepsilon^2)$, the map $D(a + b\varepsilon) = b\varepsilon$ is a nonzero derivation. The nilpotent element changes the algebraic hypothesis.

Example 15.2 (The precise defect of a forward difference)**.** On a cyclic grid, let $D_h f_j = (f_{j+1} - f_j)/h$, with $h \neq 0$. Expanding the product gives

$$D_h(fg)_j - (D_h f)_j g_j - f_j(D_h g)_j = h(D_h f)_j(D_h g)_j.$$

Thus the ordinary Leibniz identity fails unless the displayed product vanishes. A small defect on sufficiently regular sampled functions is compatible with the no-go theorem, because that theorem concerns an exact identity for every vector at fixed finite resolution.

15.1 Exact differentiation in the full matrix algebra

Move now to $M_N(\mathbb{C})$. To avoid confusing a Hermitian observable with the skew-adjoint state generator K, denote a Hermitian matrix by $B = B^*$ and define

$$\delta_B X = i[B, X] = i(BX - XB).$$

Then direct multiplication proves

$$\delta_B(XY) = (\delta_B X)Y + X\delta_B Y, \tag{42}$$
$$(\delta_B X)^* = \delta_B(X^*), \tag{43}$$
$$\mathrm{Tr}((\delta_B X)Y) = -\mathrm{Tr}(X\delta_B Y). \tag{44}$$

For the first identity, the middle terms $-iXBY$ and $iXBY$ cancel. For the second, taking the adjoint reverses the product order and conjugates i, with $B^* = B$ restoring the stated sign. For the third, cyclicity moves the leading B in $\mathrm{Tr}(BXY)$ to the end. The trace pairing is bilinear here. Combining (43) and (44) also gives the Hilbert–Schmidt identity

$$\langle \delta_B X, Y\rangle_{\mathrm{HS}} = -\langle X, \delta_B Y\rangle_{\mathrm{HS}}, \qquad \langle X, Y\rangle_{\mathrm{HS}} = \mathrm{Tr}(X^*Y).$$

These identities preserve multiplication order. They do not turn δ_B into a derivation from diagonal matrices to diagonal matrices. Usually it leaves that commutative subalgebra.

Example 15.3 (A nonzero derivative leaves the pointwise subalgebra)**.** For

$$B = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad X = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix},$$

one obtains

$$\delta_B X = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \neq 0.$$

The output is off diagonal. This is exactly how the matrix algebra permits a nonzero derivation without violating the pointwise theorem.

For Hermitian generators B_i, B_j, put $\mathcal{F}_{ij} = i[B_i, B_j]$. This curvature matrix is Hermitian, and the Jacobi identity gives

$$[\delta_{B_i}, \delta_{B_j}]X = -[[B_i, B_j], X] = i[\mathcal{F}_{ij}, X] = \delta_{\mathcal{F}_{ij}} X. \tag{45}$$

Commuting generators produce commuting derivations. A nonzero curvature matrix need not act nontrivially on every observable: its action on a particular X vanishes precisely when that X commutes with it. This matrix identity provides an algebraic curvature mechanism. A physical geometric interpretation requires the additional structures developed in the gauge and gravity papers.

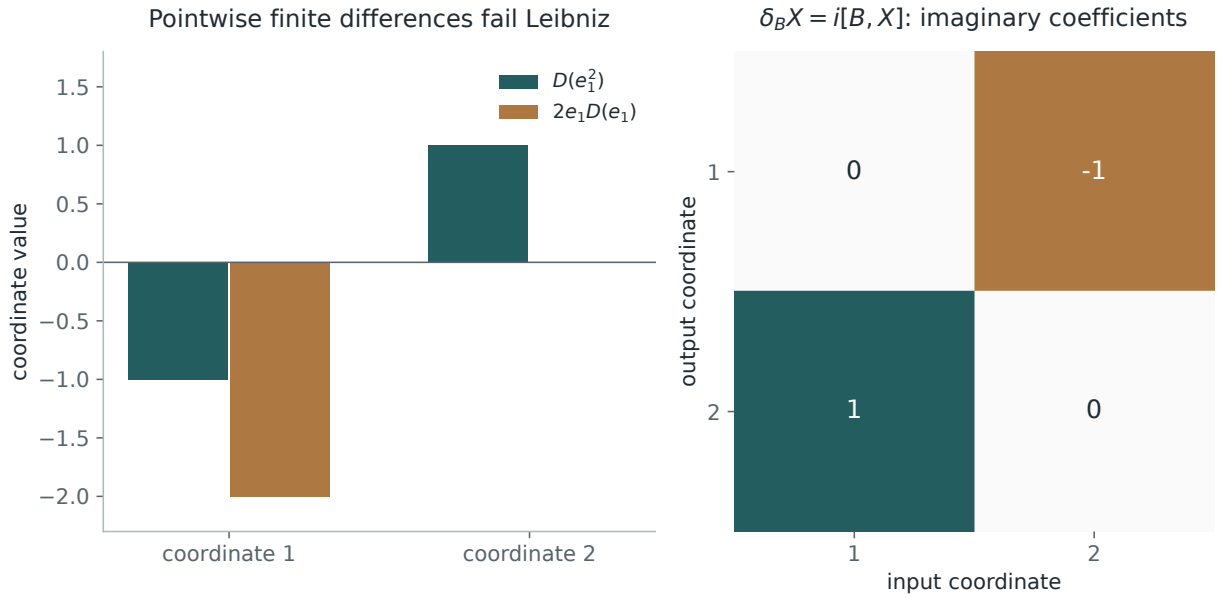


Figure 10. A finite pointwise difference has a Leibniz defect. Commutation inside a full matrix algebra obeys the Leibniz rule exactly, with its output allowed to leave the diagonal subalgebra.

15.2 Differentiated corner paths and coarse reconstruction

Let (P_a) be a finite sharp atlas. Writing $X_{ba} = P_b X P_a$ and $B_{ba} = P_b B P_a$, insertion of $I = \sum_c P_c$ gives

$$P_b \delta_B(X) P_a = i \sum_c \left(B_{bc} X_{ca} - X_{bc} B_{ca} \right). \tag{46}$$

Each term has a genuine intermediate packet c. One ordering applies X then B, while the other applies B then X. Replacing a matrix product by an entrywise product would destroy this path composition and the cancellation behind the Leibniz rule.

Summing the differentiated corners reconstructs $\delta_B X$. For a coarse partition with $Q_\alpha = \sum_{a \in C_\alpha} P_a$,

$$Q_\beta \delta_B(X) Q_\alpha = \sum_{\substack{b \in C_\beta \\ a \in C_\alpha}} P_b \delta_B(X) P_a.$$

This is exact aggregation of the already differentiated operator. It does not claim that discarding all internal matrix structure produces a new pointwise algebra with a nonzero exact derivation. The retained corners are the structure that makes the finite differential calculus work.

16 From packet calculus to field dynamics

The construction developed here has several exact outputs. It reconstructs operators from their transfer blocks. It converts skew-adjoint state dynamics into antisymmetric packet currents. It carries finite balance equations through partitions and moving frames. It identifies the additional information required by mixed states and open dynamics. It also gives precise versions of reconstruction beyond a finite atlas, where topology and domain membership become part of the statement.

The distinction between a supplied evolution and an existence theorem matters throughout the series. Differentiating a trajectory that already satisfies a state equation establishes the associated balance law. Constructing an evolution from a generator requires separate hypotheses and an existence argument. Identifying the generator of a previously constructed group is another specific task. Keeping these steps separate allows a later paper to reuse the balance identities without silently importing an unproved dynamics theorem.

The free Dirac monograph gives one concrete continuation. Fourier modes resolve the spinor field. Each mode carries a finite matrix, and a unitary Fourier map reconstructs the physical field. Finite packet consistency extends by Hilbert-space stability, while the maximal operator domain is recovered from weighted square summability. The Fourier construction provides the self-adjoint physical Hamiltonian. The following generator chapter proves its spectral-evolution identification at the level of mathematical analysis. Formal certification of that identification is recorded separately from the operator construction.

Gauge covariance and geometric refinement then extend the transport setting. Its task is to carry ordered transport through edge subdivision and cubical reconstruction while controlling the resulting defects. The moving-frame and transfer-block formulas in this volume supply part of its language. A regular prescribed connection, a coupled field evolution, and a constrained dynamical gauge theory each require their own precise mathematical setting.

The remaining parts of this book develop adic refinement, constrained gravity, and cosmological transport. Their use of the same resolution language does not eliminate their separate analytic and physical obligations. It makes those obligations comparable: specify the state space, identify the maps between resolutions, state the conserved or transported quantity, and exhibit the residual that prevents a stronger conclusion.

> **A reusable reading discipline**
>
> For each later result, locate the resolution, the transfer rule, the reconstruction statement, and the hypotheses that justify differentiation or passage to a limit. When a step fails, retain the explicit obstruction. It becomes part of the mathematical record and guides the next construction.

The same reading discipline applies across the book: exact algebra, analytic reconstruction, and physical interpretation must each carry their own hypotheses.

Sources and further reading

Project sources

[P1] Komorebi P. *Projector Refinement Physics, Volume I*. The foundational source studies cover transfer calculus, mixed and open dynamics, countable resolutions, unbounded domains, positive effects, and matrix-corner derivations. The technical source concordance identifies the relevant statements.

[P2] Komorebi P. *Projector Refinement Physics, Volume I: Technical Supplement*. Claim statements, formula registries, computational checks, and the hypotheses delimiting their conclusions.

[P3] Komorebi P. *Komorebi Physics formal source*. The foundations verification appendix gives declaration identities, retained compiler evidence, and the limits of its coverage.

[P4] Komorebi P. *Fourier Packets and the Continuum Dirac Operator*. The free-field construction in this book.

Classical context

[C1] G. Lindblad. "On the generators of quantum dynamical semigroups." *Communications in Mathematical Physics* **48** (1976), 119–130. https://doi.org/10.1007/BF01608499.

[C2] V. Gorini, A. Kossakowski, and E. C. G. Sudarshan. "Completely positive dynamical semigroups of N-level systems." *Journal of Mathematical Physics* **17**, no.5 (1976), 821–825. https://doi.org/10.1063/1.522979.

[C3] John Watrous. *Theory of Quantum Information*, CS766/QIC820, Fall2011, Lecture5, Section5.1, Theorem5.1. The finite-dimensional Naimark construction is given by the square-root isometry and coordinate projectors. https://cs.uwaterloo.ca/~watrous/TQI-notes/TQI-notes.05.pdf.

The first two classical references identify the standard setting for the finite Lindblad expression used in this volume. Their general classification results are not premises of the elementary packet calculations proved here. The Watrous notes give a direct comparison for the finite dilation construction. The source-check record distinguishes verified bibliographic metadata, the inspected dilation proof, and the inaccessible full AIP article.

This bibliography provides context and traceability. It is not an exhaustive literature review or a claim of historical priority. The mathematical review concerns the displayed arguments and the explicitly recorded source correspondence.

Fourier Packets and the Continuum Dirac Operator

From finite spinor dynamics to a self adjoint Hamiltonian on the three torus

From finite spinors to a continuum Hamiltonian

The preceding calculus becomes a concrete field construction when Fourier modes are used as packets. Each mode carries a four-component spinor and an explicit Hermitian matrix. The field is reconstructed through an isometric Fourier map. Exact norm preservation, completion, approximation, and control of the adjoint domain justify the passage from finite calculations to continuum analysis.

The setting is a free spinor on a fixed flat three-torus. For every real mass parameter, the resulting Hamiltonian is densely defined, self-adjoint, and closed. Its domain is exactly the first-order weighted Fourier square-summability domain. These assertions include the operator domain, not merely a formal differential expression.

The Hamiltonian and the modewise unitary evolution are first constructed separately. The following chapter, *The Dirac Group and Its Generator*, connects them by proving strong continuity, equality of the full generator domain with the Hamiltonian domain, and the spectral-exponential identity. It also develops graph-domain and strong-resolvent convergence for the centered Fourier family and a quadratic spacing estimate under third-order Fourier regularity.

> **The role of this construction**
>
> Finite Clifford identities determine each mode. Hilbert-space completion reconstructs the field. Adjoint-domain equality determines the self-adjoint operator. The generator argument then identifies its evolution. These are distinct steps in one continuous mathematical construction.

The operator construction, full generator-domain identification, centered graph and strong-resolvent limits, and third-order moment estimates have recorded Lean verification. The abstract spectral-functional-calculus identity remains a written deduction without a separate formal API linkage. The technical appendices identify the exact formal statements and evidence. None of these mathematical results by itself identifies an observed particle or supplies a dynamical spacetime geometry.

Reading routes and evidence

Readers seeking the construction can begin with the conventions and Clifford matrices, follow the finite-packet estimates, and continue through coefficient completion to the adjoint calculation. Readers already comfortable with free Dirac matrices can begin at the full coefficient Hilbert space. The exact three-mode example gives a concrete state that can be followed through both continuum and centered evolution.

There are three kinds of statement in this volume. A *checked correspondence* restates mathematics represented by named declarations in the retained formal source. An *ordinary mathematical deduction* has its proof in the text but is not automatically assigned to a Lean declaration. An *illustration* is an explicitly bounded Python evaluation of a formula. The figure record specifies its parameters and numerical checks.

The finite-packet propositions retain the identifiers PH01 to PH12 from the earlier traditional paper. This preserves their correspondence with its worked example and source notes. The later sections use numbered theorems for the completed-space and unbounded-operator results. A correspondence annotation identifies the checked source layer without asserting that the prose proof is a literal transcription of a tactic script.

Object	Meaning in this volume
Mode Hamiltonian	An everywhere-defined Hermitian matrix on a four-dimensional spinor space.
Coefficient evolution	A modewise unitary map on the complete coefficient Hilbert space.
Physical evolution	The conjugate of that map under the checked Fourier equivalence.
Continuum Hamiltonian	An unbounded partial linear operator with its domain carried explicitly.
Centered approximation	The multiplier with kinetic symbol $\sin(hp_j)/h$ on the same integer mode index set.
Generator identification	The theorem in the following chapter connecting the full-domain derivative and spectral exponential to the constructed Hamiltonian.

The centered approximation is defined on continuum torus functions. A spatially sampled finite grid would require its own sampling, aliasing, and embedding conventions. Similarly, first-order regularity is expressed here by the precise Fourier condition that the sources prove. An identification with every possible weak-derivative or Sobolev-library convention would require a further equivalence theorem.

No finite count of successful numerical evaluations replaces the universal statements in the manuscript or their formal counterparts. Conversely, checking a mathematical model does not establish that the model describes observed matter. The physical setting here is a free spinor on a fixed flat periodic space, in natural units.

Resolution and finite dynamics

1 Packets and transfer blocks

The useful starting point is a resolved operator rather than a preferred coordinate matrix. Let V be a complex vector space with a finite family of projectors $(P_a)_{a\in A}$ satisfying

$$P_aP_b = \delta_{ab}P_a, \qquad \sum_{a\in A} P_a = I.$$

The component $P_a v$ is the packet of v selected by a. For a linear map $L : V \to V$, its transfer block from packet a to packet b is

$$L_{ba} = P_bLP_a.$$

Completeness gives $L = \sum_{a,b} L_{ba}$. Both indices matter. Keeping only P_aLP_a would remove every transfer between distinct packets.

For two operators L and M, insert the resolution between them:

$$P_cMLP_a = \sum_{b\in A}(P_cMP_b)(P_bLP_a).$$

This is block multiplication written without a choice of basis inside each packet. The intermediate index records the interface through which a composition passes. It is also the reason that a change of resolution must specify how fine packets assemble into coarse packets.

Suppose $P_A = \sum_{a\in C(A)} P_a$ and $P_B = \sum_{b\in C(B)} P_b$. Then

$$P_BLP_A = \sum_{b\in C(B)} \sum_{a\in C(A)} P_bLP_a.$$

The identity follows by distributing two finite sums. It says exactly which fine transfers a coarse observation retains. A scalar summary or a diagonal-only summary may retain less information and therefore needs its own loss statement.

These finite identities provide the notation used throughout the wider collection. The free Fourier model has an additional property: translation invariance makes its Hamiltonian diagonal in the mode label. Each diagonal block is still a nontrivial four-spinor matrix. The simplification belongs to this model, not to every operator that can be expressed in packets.

1.1 What changes at completion

The finite identity $\sum_a P_a = I$ becomes an approximation statement for an infinite orthogonal family. Finite partial sums must converge in the chosen topology. On the coefficient Hilbert space, canonical finite truncations converge in norm. For an unbounded Hamiltonian, approximating the state alone is insufficient to control its image. The graph approximation later in the volume controls both quantities at once.

This distinction explains the structure of the proof. Exact finite algebra supplies local formulas. Hilbert stability controls their extension to all square-summable fields. Weighted regularity specifies which fields support an unbounded operator action. Each step carries a different mathematical obligation.

2 Conventions and finite packet space

We use natural units $\hbar = c = 1$. Space is the flat unit torus

$$\mathbb{T}^3 = (\mathbb{R}/\mathbb{Z})^3,$$

equipped with normalized Haar measure dx, so $\int_{\mathbb{T}^3} 1\, dx = 1$. Integration of a periodic representative may equivalently be taken over $[0,1)^3$ with Lebesgue measure. A Fourier mode is $n = (n_1, n_2, n_3) \in \mathbb{Z}^3$, its momentum is $p(n) = 2\pi n$, and its scalar plane wave is

$$e_n(x) = e^{2\pi i n \cdot x} = e^{ip(n)\cdot x}.$$

Components of a spinor are numbered $0, 1, 2, 3$, in the exact order of the source representation. For $v, w \in \mathbb{C}^4$,

$$\langle v, w \rangle = \sum_{r=0}^{3} \overline{v_r} w_r, \qquad \|v\|_2^2 = \sum_{r=0}^{3} |v_r|^2.$$

The inner product is conjugate-linear in its first argument and linear in its second. The adjoint is therefore the conjugate transpose. Unless indicated otherwise, a spinor norm is this Euclidean norm and an operator norm is its induced operator norm. The field pairing is

$$\langle f, g \rangle_{L^2} = \int_{\mathbb{T}^3} \langle f(x), g(x) \rangle \, dx$$

on $L^2(\mathbb{T}^3; \mathbb{C}^4)$. Fix a finite set $S \subset \mathbb{Z}^3$. A packet with amplitudes $a_n \in \mathbb{C}^4$ is

$$f_a(x) = \sum_{n \in S} e_n(x) a_n.$$

The associated coefficient space is the finite Hilbert direct sum

$$\mathcal{K}_S = \bigoplus_{n \in S} \mathbb{C}^4, \qquad \|a\|_{\mathcal{K}_S}^2 = \sum_{n \in S} \|a_n\|_2^2.$$

The packet space $\mathcal{F}_S = \{f_a : a \in \mathcal{K}_S\}$ will be identified isometrically with $\mathcal{K}_S$. Its elements are smooth periodic functions, so the differential and centered-translation calculations below raise no regularity question. Empty support is allowed: both spaces then contain only zero and all packet identities reduce to $0 = 0$.

The parameters m, h, t are arbitrary real numbers. The mathematics does not require $m \geq 0$; the usual physical rest-mass interpretation uses that additional convention. The shell value E is always a nonnegative square root, even when $m < 0$. Mesh limits in the checked statements are $h \to 0^+$. Some explicit estimates proved here hold for every real h, since the centered symbol is even in h.

3 The coordinate Clifford representation

The spatial generators and mass generator are

$$\alpha_x = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix}, \qquad \alpha_y = \begin{pmatrix} 0 & 0 & 0 & -i \\ 0 & 0 & i & 0 \\ 0 & -i & 0 & 0 \\ i & 0 & 0 & 0 \end{pmatrix},$$

$$\alpha_z = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \\ 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \end{pmatrix}, \qquad \beta = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}.$$

In particular, $\alpha_y v = (-iv_3, iv_2, -iv_1, iv_0)$; this fixes a sign that cannot be changed silently in a worked example. The matrices act on column spinors. Grouping the first two and last two components gives their block form without any coordinate permutation:

$$\alpha_j = \begin{pmatrix} 0 & \sigma_j \\ \sigma_j & 0 \end{pmatrix}, \qquad \beta = \begin{pmatrix} I_2 & 0 \\ 0 & -I_2 \end{pmatrix},$$

where

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

Proposition PH01 (Hermitian Clifford relations). The four displayed matrices are Hermitian involutions. Distinct spatial generators anticommute, and each spatial generator anticommutes with β:

$$\alpha_j^* = \alpha_j, \quad \beta^* = \beta, \quad \alpha_j^2 = \beta^2 = I_4,$$

$$\alpha_j\alpha_k + \alpha_k\alpha_j = 0 \quad (j \neq k), \qquad \alpha_j\beta + \beta\alpha_j = 0.$$

Proof. The real diagonal entries of β are $1, 1, -1, -1$, so its adjoint and square have the stated forms. Each σ_j is Hermitian and squares to I_2, by direct multiplication. Their products are

$$\sigma_x\sigma_y = i\sigma_z, \qquad \sigma_y\sigma_z = i\sigma_x, \qquad \sigma_z\sigma_x = i\sigma_y,$$

with reversed products equal to their negatives. For example,

$$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}\begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} = \begin{pmatrix} i & 0 \\ 0 & -i \end{pmatrix} = i\sigma_z.$$

The other two follow by the same two-by-two multiplication; reversal follows by taking adjoints. The block form gives

$$\alpha_j^* = \begin{pmatrix} 0 & \sigma_j^* \\ \sigma_j^* & 0 \end{pmatrix} = \alpha_j, \qquad \alpha_j\alpha_k = \begin{pmatrix} \sigma_j\sigma_k & 0 \\ 0 & \sigma_j\sigma_k \end{pmatrix}.$$

Thus $\alpha_j^2 = I_4$ and the spatial anticommutators vanish. Finally,

$$\alpha_j\beta = \begin{pmatrix} 0 & -\sigma_j \\ \sigma_j & 0 \end{pmatrix}, \qquad \beta\alpha_j = \begin{pmatrix} 0 & \sigma_j \\ -\sigma_j & 0 \end{pmatrix},$$

whose sum is zero. Every identity is an identity of the exact source matrices. □

4 Continuum and centered mode Hamiltonians

For $p \in \mathbb{R}^3$, define

$$H(m,p) = m\beta + p_x\alpha_x + p_y\alpha_y + p_z\alpha_z.$$

Its coordinates are

$$H(m,p) = \begin{pmatrix} m & 0 & p_z & p_x - ip_y \\ 0 & m & p_x + ip_y & -p_z \\ p_z & p_x - ip_y & -m & 0 \\ p_x + ip_y & -p_z & 0 & -m \end{pmatrix}.$$

On mode n, the continuum Hamiltonian is $H_n = H(m, p(n))$. It is the symbol of the differential expression

$$-i\sum_{j=x,y,z} \alpha_j\partial_j + m\beta$$

on the finite packet space: $\partial_j e_n = ip_j(n)e_n$, so applying the expression to $e_n a_n$ gives $e_n H_n a_n$. For real spacing h, set

$$q_j(h,p) = \begin{cases} \sin(hp_j)/h, & h \neq 0, \\ p_j, & h = 0. \end{cases}$$

Write $q_h(n) = q(h, p(n))$ and $H_{h,n} = H(m, q_h(n))$. The function q_j is even in h, and its value at zero is justified by $\sin z/z \to 1$. For $h \neq 0$, a centered translation difference on a periodic smooth function is

$$(D_{h,j}f)(x) = \frac{f(x + he_j) - f(x - he_j)}{2h}.$$

Here e_j is the j-th spatial unit vector, distinct from the Fourier characters e_n and spinor vectors. Its action on a plane wave is exactly

$$D_{h,j}e_n = e_n\frac{e^{ihp_j(n)} - e^{-ihp_j(n)}}{2h} = i\frac{\sin(hp_j(n))}{h}e_n.$$

The centered differential expression $-i\sum_j \alpha_j D_{h,j} + m\beta$ therefore has symbol $H_{h,n}$. Its signs follow from $-i \cdot i = 1$. At $h = 0$, the symbol equals H_n by definition. These formulas take place on the continuum torus and use no sampled spatial data.

Proposition PH02 (Hermiticity and the mass shell). For every real m and $p \in \mathbb{R}^3$,

$$H(m,p)^* = H(m,p), \qquad H(m,p)^2 = (m^2 + |p|^2)I_4.$$

In particular, $H_n, H_{h,n}$ are Hermitian for all real m, h, and

$$H_n^2 = E_n^2 I_4, \quad E_n = \sqrt{m^2 + |p(n)|^2},$$

$$H_{h,n}^2 = E_{h,n}^2 I_4, \quad E_{h,n} = \sqrt{m^2 + |q_h(n)|^2}.$$

Proof. Adjoint is additive, and a real multiple of a Hermitian matrix is Hermitian. PH01 therefore proves the first identity. Expanding the square gives the square terms $m^2\beta^2 + \sum_j p_j^2\alpha_j^2$, the terms $mp_j(\beta\alpha_j + \alpha_j\beta)$, and the terms $p_jp_k(\alpha_j\alpha_k + \alpha_k\alpha_j)$ for $j < k$. The anticommutators vanish and all four squares are I_4, giving $(m^2 + \sum_j p_j^2)I_4$. Substitution proves both mode statements. □

Hermiticity here is a statement about matrices acting on all of $\mathbb{C}^4$. Every linear map on this finite space is bounded and everywhere defined. This is not a construction of the adjoint or domain of an unbounded differential operator on the full field space.

5 Spectral projectors and exact mode evolution

Let $H = H(m,p)$ and $E = \sqrt{m^2 + |p|^2}$. Define the exact propagator

$$U_H(t) = e^{-itH} = \sum_{k=0}^{\infty} \frac{(-itH)^k}{k!}.$$

The series converges absolutely in operator norm, being dominated by $\sum_k (|t|\|H\|)^k/k!$. This is exact continuous-time evolution, with no time step or series truncation.

Proposition PH03 (spectral splitting and closed propagator). If $E > 0$, define

$$P_+ = \frac{1}{2}\left(I_4 + \frac{H}{E}\right), \qquad P_- = \frac{1}{2}\left(I_4 - \frac{H}{E}\right).$$

Then $P_\pm^* = P_\pm$, $P_\pm^2 = P_\pm$, $P_+P_- = P_-P_+ = 0$, $P_+ + P_- = I_4$, and $HP_\pm = \pm EP_\pm$. Both projectors have rank two, and

$$U_H(t) = e^{-itE}P_+ + e^{itE}P_- = \cos(tE)I_4 - i\frac{\sin(tE)}{E}H.$$

If $E = 0$, then $H = 0$ and $U_H(t) = I_4$. Equivalently, the closed expression is valid with the continuous convention $\sin(tE)/E = t$ at $E = 0$.

Proof. For $E > 0$, the Hermitian matrix $K = H/E$ satisfies $K^2 = I_4$. Thus

$$P_\pm^2 = \tfrac{1}{4}(I_4 \pm 2K + K^2) = \tfrac{1}{2}(I_4 \pm K) = P_\pm,$$

and $P_+P_- = \frac{1}{4}(I_4 - K^2) = 0$; the reversed product is the same. Their adjoints equal themselves and their sum is I_4. Also $K(I_4 \pm K) = K \pm I_4 = \pm(I_4 \pm K)$, proving $HP_\pm = \pm EP_\pm$. Each displayed Clifford matrix has trace zero. Hence tr $H = 0$ and tr $P_\pm = 2$. An idempotent splits the space into its range and kernel, acting as identity and zero respectively. Its trace is therefore its rank, proving rank two.

The identity $I_4 = P_+ + P_-$ and the two eigenvalue identities give, by induction, $H^k = E^k P_+ + (-E)^k P_-$ for every $k \geq 0$. Substitute into the absolutely convergent series to obtain the spectral expression. Expanding its projectors gives the cosine and sine expression.

When $E = 0$, the sum $m^2 + |p|^2$ of nonnegative squares is zero. Therefore $m = 0, p = 0$, and $H = 0$. Its exponential is I_4. The quotient $\sin(tE)/E$ tends to t at zero; multiplying its filled value by $H = 0$ gives the same identity. No division by zero is used. □

The proof applies unchanged to $H_{h,n}$, with energy $E_{h,n}$. A positive-energy continuum mode can have $E_{h,n} = 0$ at a nonzero spacing when $m = 0$ and all centered components vanish. That mode uses the zero-energy branch. Positivity of the continuum energy does not license division by the centered energy at every spacing.

Proposition PH04 (group law, unitarity, and stability). For real s, t,

$$U_H(0) = I_4, \qquad U_H(s+t) = U_H(s)U_H(t),$$

$$U_H(-t) = U_H(t)^{-1} = U_H(t)^*, \qquad U_H(t)^*U_H(t) = U_H(t)U_H(t)^* = I_4.$$

Consequently,

$$\langle U_H(t)v, U_H(t)w\rangle = \langle v, w\rangle, \qquad \|U_H(t)v\|_2 = \|v\|_2$$

for all spinors. These statements hold for both mode Hamiltonians at every real m, h, t.

Proof. When $E > 0$, the spectral form and orthogonality of $P_\pm$ give

$$U_H(s)U_H(t) = e^{-i(s+t)E}P_+ + e^{i(s+t)E}P_- = U_H(s+t).$$

At zero time the sum is I_4. Hermiticity of the projectors and reality of tE give $U_H(t)^* = e^{itE}P_+ + e^{-itE}P_- = U_H(-t)$. Setting $s = -t$ proves the inverse identities. When $E = 0$, every assertion follows from $U_H(t) = I_4$.

For any matrix A, expansion and interchange of two finite sums gives $\langle Av, w\rangle = \langle v, A^*w\rangle$. Thus

$$\langle U_H(t)v, U_H(t)w\rangle = \langle v, U_H(t)^*U_H(t)w\rangle = \langle v, w\rangle.$$

Take $w = v$ and then nonnegative square roots to obtain norm preservation. PH02 supplies the mode cases. □

An equivalent bounded-operator argument is represented in the formal source: $-itH$ is skew-adjoint, and its exponential is unitary. Our spectral proof supplies additional explicit projectors and formulas.

6 Fourier orthogonality and packet conservation

Let $J_n : \mathbb{C}^4 \to L^2(\mathbb{T}^3;\mathbb{C}^4)$ be $J_n v = e_n v$.

Proposition PH05 (mode isometry and finite Parseval identity). For modes n, k and spinors v, w,

$$\langle J_n v, J_k w\rangle_{L^2} = \begin{cases} \langle v, w\rangle, & n = k, \\ 0, & n \neq k. \end{cases}$$

In particular, J_n is a linear isometry and

$$\left\| \sum_{n\in S} J_n a_n \right\|_{L^2}^2 = \sum_{n\in S} \|a_n\|_2^2.$$

Proof. The constant spinors factor out of the scalar integral:

$$\langle J_n v, J_k w\rangle_{L^2} = \langle v, w\rangle \int_{\mathbb{T}^3} e^{2\pi i(k-n)\cdot x}\, dx.$$

The integral is the product of three one-dimensional integrals. For integer r, the integral of $e^{2\pi i r x}$ over $[0,1)$ is 1 when $r = 0$; for $r \neq 0$, it is $(e^{2\pi i r} - 1)/(2\pi i r) = 0$. The product is therefore 1 exactly when $n = k$. Linearity is immediate, and setting $k = n, w = v$ proves isometry. Expansion of the squared norm of the finite sum gives a double sum of pairings. The off-diagonal pairings vanish and the diagonal pairings are $\|a_n\|_2^2$. □

Define the evolved packets

$$f(t,x) = \sum_{n\in S} e_n(x) U_{H_n}(t) a_n, \qquad f_h(t,x) = \sum_{n\in S} e_n(x) U_{H_{h,n}}(t) a_n.$$

Both remain in $\mathcal{F}_S$. Their operators $U_S(t), U_{h,S}(t)$ have the finite block diagonal coefficient actions with blocks $U_{H_n}(t), U_{H_{h,n}}(t)$.

Theorem PH06 (exact finite-packet norm conservation). For every real m, h, t, finite S, and coefficient family a,

$$\|f(t)\|_{L^2}^2 = \|f_h(t)\|_{L^2}^2 = \sum_{n\in S} \|a_n\|_2^2 = \|f_a\|_{L^2}^2.$$

Both packet evolutions preserve the L^2 inner product between two packets in $\mathcal{F}_S$, and each forms a one-parameter unitary group there.

Proof. Apply PH05 to the evolved coefficients and PH04 on each mode:

$$\|f_h(t)\|_{L^2}^2 = \sum_{n\in S} \|U_{H_{h,n}}(t) a_n\|_2^2 = \sum_{n\in S} \|a_n\|_2^2.$$

The continuum argument is identical. For coefficient families a, b, the same two propositions give

$$\langle f_{h,a}(t), f_{h,b}(t)\rangle_{L^2} = \sum_{n\in S} \langle U_{H_{h,n}}(t) a_n, U_{H_{h,n}}(t) b_n\rangle = \sum_{n\in S} \langle a_n, b_n\rangle.$$

Composition of coefficient blocks proves the packet group law, and inverse time proves surjectivity. This establishes unitarity on $\mathcal{F}_S$ using only finite sums. □

The quantity conserved is squared Hilbert norm, or total probability for a normalized state. The parameter m, the positive shell value E_n, and Hamiltonian expectation are different quantities. Calling squared norm a packet mass does not identify it with the rest mass or physical Hamiltonian energy.

7 Generator consistency and dispersion error

The mode Hamiltonian difference contains no mass term:

$$\Delta_{h,n} := H_{h,n} - H_n = \sum_{j=x,y,z} \delta_j(h,n)\alpha_j, \qquad \delta_j(h,n) = q_j(h,p(n)) - p_j(n).$$

Define

$$C(n) = \frac{1}{6}\left(\sum_{j=x,y,z} |p_j(n)|^6\right)^{1/2}.$$

This constant depends on the fixed mode and increases cubically with momentum scale. That dependence is essential in a fixed-cutoff result.

Proposition PH07 (quadratic generator consistency). For every real h and fixed mode n,

$$\|\Delta_{h,n}\| = |\delta(h,n)| \le h^2 C(n).$$

For every fixed finite packet,

$$\left\|\sum_{n\in S} e_n \Delta_{h,n} a_n\right\|_{L^2}^2 = \sum_{n\in S} |\delta(h,n)|^2 \|a_n\|_2^2,$$

and

$$\left\|\sum_{n\in S} e_n \Delta_{h,n} a_n\right\|_{L^2} \le h^2 \left(\sum_{n\in S} C(n)^2 \|a_n\|_2^2\right)^{1/2}.$$

Proof. First prove the scalar estimate without a sign restriction on h. The fundamental theorem of calculus gives $|\sin u| \le |u|$, since $\sin u = \int_0^u \cos v\, dv$ and $|\cos v| \le 1$. Therefore

$$|1 - \cos u| = 2\sin^2(u/2) \le u^2/2.$$

For $z \ge 0$, integrate this bound in $\sin z - z = \int_0^z (\cos u - 1)\, du$ to obtain $|\sin z - z| \le z^3/6$. Oddness of the residual proves $|\sin z - z| \le |z|^3/6$ for negative z. For $h \ne 0$, substitute $z = hp_j(n)$ and divide by $|h|$:

$$|\delta_j(h,n)| \le \frac{h^2 |p_j(n)|^3}{6}.$$

For $h = 0$, the residual is zero by definition and the bound still holds. Squaring and adding proves $|\delta(h, n)| \le h^2 C(n)$.

The Hermitian matrix $\Delta_{h,n}$ satisfies $\Delta_{h,n}^2 = |\delta(h, n)|^2 I_4$ by PH01. Hence

$$\|\Delta_{h,n} v\|_2^2 = \langle v, \Delta_{h,n}^* \Delta_{h,n} v\rangle = |\delta(h, n)|^2 \|v\|_2^2.$$

Its induced operator norm is exactly $|\delta(h, n)|$: the upper bound holds on every vector and equality holds on every unit vector. Apply PH05 to the residual coefficients for the packet equality, substitute the scalar bounds, and take nonnegative square roots. □

The checked generator results state corresponding $O(h^2)$ estimates as $h \to 0^+$. The explicit constant and all-real inequality above are traditional deductions. Precisely, $r(h) = O(h^2)$ means that some $M, \varepsilon > 0$ satisfy $\|r(h)\| \le Mh^2$ for $0 < h < \varepsilon$. Fixed support and amplitudes permit finite summation of such bounds; no estimate uniform over all integer modes follows.

The reverse triangle inequality in $\mathbb{R}^4$, applied to $(m, q_h(n))$ and $(m, p(n))$, gives

$$|E_{h,n} - E_n| \le |q_h(n) - p(n)| \le h^2 C(n).$$

Moreover, $|\sin(hp_j)/h| \le |p_j|$ for $h \neq 0$, and the filled value is equal at zero. Thus $0 \le E_{h,n} \le E_n$. For a fixed mode, Taylor's formula gives

$$q_j(h, p(n)) = p_j(n) - \frac{h^2 p_j(n)^3}{6} + O(h^4),$$

$$E_{h,n}^2 = E_n^2 - \frac{h^2}{3} \sum_j p_j(n)^4 + O(h^4).$$

For clarity, Taylor's theorem through degree four has zero even coefficients for sine, and its fifth-derivative remainder is at most $|hp_j|^5/120$. Division by $|h|$ gives $O(h^4)$; squaring the three expansions and summing preserves that order for a fixed mode. If $E_n > 0$, the square root is smooth near E_n^2, so

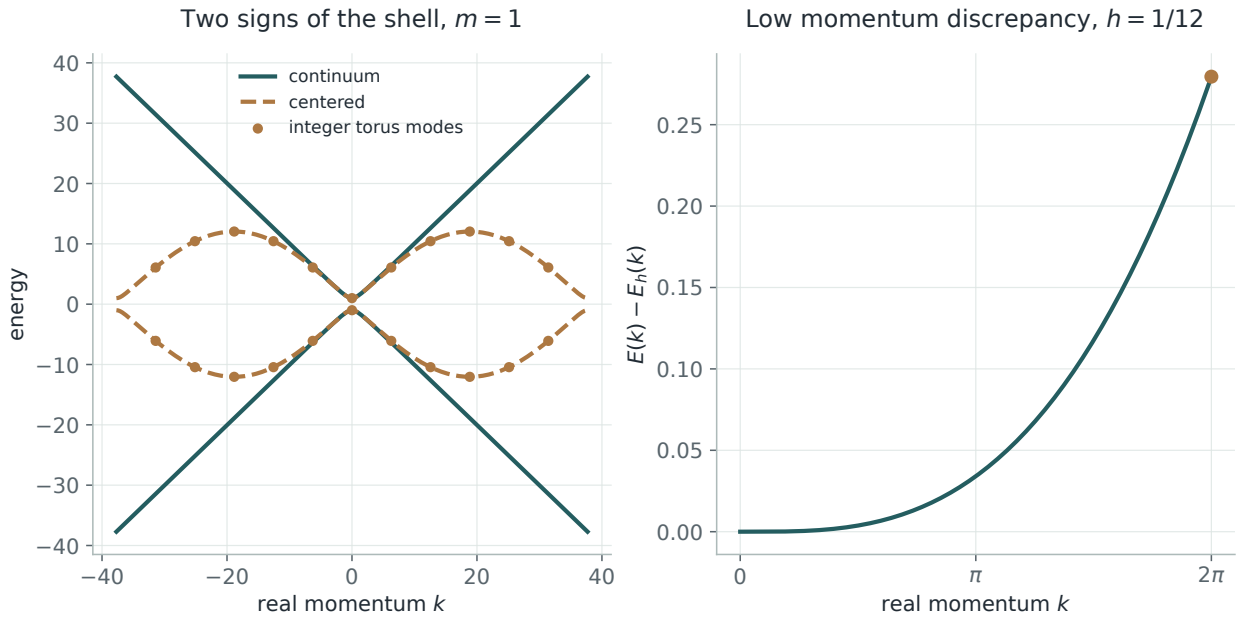
$$E_{h,n} = E_n - \frac{h^2}{6E_n} \sum_j p_j(n)^4 + O(h^4).$$

If $E_n = 0$, then $m = 0, n = 0$ and $E_{h,n} = 0$ exactly for every h. The expansion containing $1/E_n$ is not used.

For fixed time, the phase error on either energy branch obeys

$$|e^{\pm itE_{h,n}} - e^{\pm itE_n}| \le |t|\,|E_{h,n} - E_n|,$$

by integrating the derivative of $s \mapsto e^{\pm its}$ between the two real energies. Phase comparison alone does not compare full spinor evolution: the eigenspaces also depend on the kinetic vector. The operator argument below controls both effects.



The continuum and centered shell branches along one momentum axis, with $m = 1$, $h = 1/12$. Marked points are integer torus modes; intermediate points use the real-symbol extension.

8 Fixed-time propagator consistency

Theorem PH08 (fixed-mode and fixed-packet convergence). For fixed real m, t and fixed n,

$$U_{H_{h,n}}(t) - U_{H_n}(t) = O(h^2) \quad (h \to 0^+).$$

For fixed finite S and amplitudes a,

$$f_h(t) - f(t) = O(h^2) \quad \text{in } L^2(\mathbb{T}^3; \mathbb{C}^4),$$

and in particular $\|f_h(t) - f(t)\|_{L^2} \to 0$. The exact residual identity is

$$\|f_h(t) - f(t)\|_{L^2}^2 = \sum_{n \in S} \|(U_{H_{h,n}}(t) - U_{H_n}(t))a_n\|_2^2.$$

Proof. Give a fixed-time proof independent of the later compact-time extension. For arbitrary matrices X, Y, telescoping products give

$$X^k - Y^k = \sum_{r=0}^{k-1} X^{k-1-r}(X - Y)Y^r \quad (k \geq 1).$$

No commutation is assumed. Expanding the summands gives successive differences $X^{k-r}Y^r - X^{k-1-r}Y^{r+1}$, which cancel. If $\|X\|, \|Y\| \leq R$, then

$$\|X^k - Y^k\| \leq kR^{k-1}\|X - Y\|.$$

Subtract the absolutely convergent exponential series and sum:

$$\|e^X - e^Y\| \le \|X - Y\| \sum_{k=1}^{\infty} \frac{kR^{k-1}}{k!} = e^R \|X - Y\|.$$

PH02 gives $\|H_n\| = E_n, \|H_{h,n}\| = E_{h,n}$, by the norm-squared argument of PH07. Since $E_{h,n} \le E_n$, take $X = -itH_{h,n}$, $Y = -itH_n$, and $R = |t|E_n$. The result is

$$\|U_{H_{h,n}}(t) - U_{H_n}(t)\| \le |t| e^{|t|E_n} h^2 C(n).$$

This proves the modewise order, including $t = 0$. The difference of packets is the finite sum of the mode residuals, so PH05 gives its exact residual identity. Thus

$$\|f_h(t) - f(t)\|_{L^2} \le h^2 \left(\sum_{n \in S} |t|^2 e^{2|t|E_n} C(n)^2 \|a_n\|_2^2 \right)^{1/2}.$$

The constant is finite for fixed t, S, a, proving the order and convergence. □

The formal algebraic-coordinate proof obtains this order from continuity and local differentiability of the Banach-algebra exponential. Its operator norm belongs to the algebraic coordinate spinor type. The coefficient conversion in its packet definitions is continuous, so its checked torus-L^2 convergence is genuine. Our Euclidean PH08 also has the direct proof above. The following bridge explains why its coordinates agree with the separately defined Hilbert propagators.

9 The algebraic-to-Hilbert coordinate bridge

Let $\mathcal{A}$ be the coordinate space $\{0, 1, 2, 3\} \to \mathbb{C}$ with maximum norm $\|v\|_\infty = \max_r |v_r|$, and let $\mathcal{H}$ be the same coordinates with the Euclidean inner product. The map

$$T : \mathcal{A} \longrightarrow \mathcal{H}, \qquad (Tv)_r = v_r$$

is complex-linear and bijective; its inverse also retains the coordinates. It is not generally an isometry between these norms. The inequalities

$$\|v\|_\infty \le \|Tv\|_2 \le 2\|v\|_\infty$$

show that both directions are continuous. The first follows because each coordinate square is bounded by the sum; the second because each of four squares is at most $\|v\|_\infty^2$.

Proposition PH09 (coordinate conjugacy of generators and exponentials). Let A be a complex-linear endomorphism of $\mathcal{A}$, and let $A^{\mathcal{H}} = TAT^{-1}$. For every complex z,

$$e^{zA^{\mathcal{H}}} = Te^{zA}T^{-1}.$$

For the source Clifford actions, the continuum and centered Hilbert generators are exactly these conjugates of their algebraic generators. Their exact propagators therefore have the same coordinates after applying T. The algebraic-coordinate packet, transported coefficientwise by T, agrees with the Hilbert packet for the corresponding amplitudes.

Proof. Induction gives $(TAT^{-1})^k = TA^kT^{-1}$. For $k = 0$, use $TT^{-1} = I$; for the induction step, cancel the adjacent $T^{-1}T$. Every exponential partial sum therefore satisfies

$$\sum_{k=0}^{N} \frac{z^k (A^{\mathcal{H}})^k}{k!} = T \left(\sum_{k=0}^{N} \frac{z^k A^k}{k!} \right) T^{-1}.$$

The exponential series converges in both operator norms. Conjugation is continuous: writing induced norms in the relevant spaces gives

$$\|TBT^{-1}\|_2 \le \|T\| \, \|B\|_\infty \, \|T^{-1}\| \le 2\|B\|_\infty.$$

Pass to the limit to obtain the identity.

The source's transported Clifford action is defined by $v \mapsto T(A(T^{-1}v))$. The same is true for the displayed diagonal beta action. Conjugation preserves real coefficients and finite sums, proving the Hamiltonian equality. Apply the exponential identity with $z = -it$. Substituting each coefficient equality into the finite Fourier sum then proves equality of the packets. □

The norm bound also transports quadratic operator errors between realizations. The bridge is a traditional proof in this paper. The checked source contains inverse coordinate identities and the pairing identification, but no complete generator-and-exponential conjugacy theorem. PH09's conjugacy itself remains an ordinary proof in this text. Blocks 44 and 45 separately establish Hilbert stability and consistency directly.

10 An exact three-mode packet

Example PH10 (axis and diagonal modes at $h = 1/12$). Set $m = 1$, $h = 1/12$, and

$$S = \{n_0, n_A, n_D\}, \qquad n_0 = (0,0,0), \quad n_A = (1,0,0), \quad n_D = (1,1,1).$$

Let $e_0 = (1,0,0,0)^T$ denote the first spinor basis vector, and choose

$$a_{n_0} = \frac{1}{\sqrt{2}} e_0, \qquad a_{n_A} = \frac{1}{2} e_0, \qquad a_{n_D} = \frac{1}{2} e_0.$$

Their squared norms are $1/2, 1/4, 1/4$, so the packet is normalized. Its continuum and centered momenta are

Mode	$p(n)$	$q_h(n)$
n_0	$(0,0,0)$	$(0,0,0)$
n_A	$(2\pi,0,0)$	$(6,0,0)$
n_D	$(2\pi,2\pi,2\pi)$	$(6,6,6)$

Each nonzero component has $hp_j = \pi/6$, so $\sin(hp_j)/h = (1/2)/(1/12) = 6$. Thus

Mode	E_n	$E_{h,n}$
n_0	1	1
n_A	$\sqrt{1+4\pi^2}$	$\sqrt{37}$
n_D	$\sqrt{1+12\pi^2}$	$\sqrt{109}$

These are exact symbolic equalities. The centered component 6 is distinct from the integer mode label 1 and continuum momentum 2π.

For any real kinetic vector r, the exact matrix gives

$$H(1,r)e_0 = (1,0,r_z,r_x+ir_y)^T.$$

Set $E_r=\sqrt{1+|r|^2}$, $c_r(t)=\cos(E_rt)$, and $s_r(t)=\sin(E_rt)/E_r$. PH03 gives

$$u_r(t)=U_{H(1,r)}(t)e_0=\begin{pmatrix} c_r(t)-is_r(t)\\ 0\\ -ir_zs_r(t)\\ (r_y-ir_x)s_r(t)\end{pmatrix}.$$

In particular, $u_0(t)=e^{-it}e_0$. For the centered axis and diagonal modes,

$$u_A(t)=\begin{pmatrix}\cos(\sqrt{37}t)-i\sin(\sqrt{37}t)/\sqrt{37}\\ 0\\ 0\\ -6i\sin(\sqrt{37}t)/\sqrt{37}\end{pmatrix},$$

$$u_D(t)=\begin{pmatrix}\cos(\sqrt{109}t)-i\sin(\sqrt{109}t)/\sqrt{109}\\ 0\\ -6i\sin(\sqrt{109}t)/\sqrt{109}\\ 6(1-i)\sin(\sqrt{109}t)/\sqrt{109}\end{pmatrix}.$$

The complete centered packet is

$$f_h(t,x)=\frac{e^{-it}}{\sqrt{2}}e_0+\frac{1}{2}e^{2\pi ix_1}u_A(t)+\frac{1}{2}e^{2\pi i(x_1+x_2+x_3)}u_D(t).$$

The continuum packet has the same form with u_A,u_D replaced by $u_{(2\pi,0,0)}$ and $u_{(2\pi,2\pi,2\pi)}$. At $t=0$, both equal

$$f_a(x)=\left(\frac{1}{\sqrt{2}}+\frac{1}{2}e^{2\pi ix_1}+\frac{1}{2}e^{2\pi i(x_1+x_2+x_3)}\right)e_0.$$

The spinor formula verifies stability directly:

$$\|u_r(t)\|_2^2=c_r(t)^2+s_r(t)^2(1+r_z^2+r_x^2+r_y^2)=\cos^2(E_rt)+\sin^2(E_rt)=1.$$

Here $|c-is|^2 = c^2+s^2$ and $|r_y - ir_x|^2 = r_x^2+r_y^2$, since the coefficients are real. Fourier orthogonality gives

$$\|f_h(t)\|_{L^2}^2 = \tfrac{1}{2}+\tfrac{1}{4}+\tfrac{1}{4} = 1, \qquad \|f(t)\|_{L^2}^2 = 1.$$

This is an integrated statement. The initial pointwise density is not constant: at $x = 0$, it is $(1 + 1/\sqrt{2})^2 = 3/2 + \sqrt{2}$. Cross terms can change local density while their torus integrals vanish. Normalization does not mean that each point has density one.

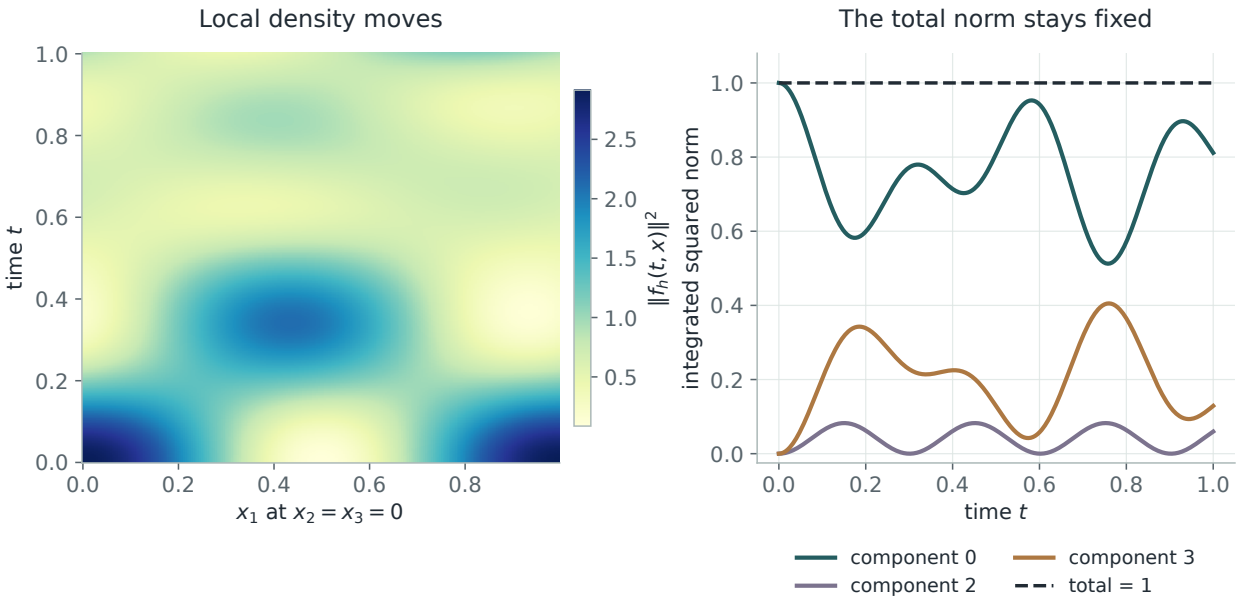
The error also has an exact expression. For fixed mode, let $r = p(n)$, $q = q_h(n)$, with the corresponding c_r, s_r, c_q, s_q. Coordinate subtraction gives

$$\|u_q(t) - u_r(t)\|_2^2 = (c_q - c_r)^2 + (s_q - s_r)^2 + \sum_{j=x,y,z} (q_j s_q - r_j s_r)^2.$$

For the selected packet, the zero mode contributes no error, so

$$\begin{aligned}\|f_h(t) - f(t)\|_{L^2}^2 = {} & \frac{1}{4}\|u_{(6,0,0)}(t) - u_{(2\pi,0,0)}(t)\|_2^2 \\ & + \frac{1}{4}\|u_{(6,6,6)}(t) - u_{(2\pi,2\pi,2\pi)}(t)\|_2^2.\end{aligned}$$

A plot or numerical evaluation illustrates this exact expression; conservation and normalization have already been proved symbolically.



The centered three-mode packet at $m = 1$, $h = 1/12$ changes local density and spinor composition while keeping total squared norm one. The image evaluates the closed formula at finitely many times.

11 A finite-matrix compact-time extension

The following proof gives an explicit Duhamel bound, including its sharp coefficient. This is an ordinary mathematical deduction with a complete proof. The formal compact-time convergence argument uses exponential continuity and is identified separately in the verification appendix.

Lemma (finite Hermitian Duhamel estimate). Let A, B be Hermitian matrices on a nonzero finite-dimensional complex Hilbert space. For every real t,

$$\|e^{-itA} - e^{-itB}\| \le |t|\,\|A - B\|.$$

The same inequality holds trivially on the zero space.

Proof. First justify differentiation. On $|s| \le R$, the series for e^{-isB} converges uniformly, and its derivative series is norm-dominated by

$$\sum_{k=1}^{\infty} \frac{kR^{k-1}\|B\|^k}{k!} = \|B\|e^{R\|B\|}.$$

Termwise differentiation therefore gives

$$\frac{d}{ds}e^{-isB} = -iBe^{-isB} = -ie^{-isB}B.$$

The same holds for A. Since A commutes with every power of itself, it commutes with e^{-iuA}. Fix t and define $F(s) = e^{-i(t-s)A}e^{-isB}$. The product rule gives

$$F'(s) = ie^{-i(t-s)A}(A - B)e^{-isB}.$$

The endpoints are $F(0) = e^{-itA}$ and $F(t) = e^{-itB}$. Integrating entrywise in the finite matrix space gives the oriented-integral identity

$$e^{-itB} - e^{-itA} = i\int_0^t e^{-i(t-s)A}(A - B)e^{-isB}\,ds.$$

For Hermitian matrices the two exponentials in the integrand are unitary. This does not rely on the special Clifford algebra: the series gives $(e^{-iuA})^* = e^{iuA}$, and the Cauchy product of the two absolutely convergent series gives $e^{iuA}e^{-iuA} = I$. Indeed, the coefficient of A^k for $k > 0$ is $(iu)^k \sum_{r=0}^{k} \binom{k}{r}(-1)^r/k! = 0$, and the constant coefficient is I. A unitary matrix on a nonzero space has norm one. Submultiplicativity therefore bounds the integrand norm by $\|A - B\|$. For $t \ge 0$, integration gives $t\|A - B\|$; for $t < 0$, reversing the interval gives $(-t)\|A - B\|$. This proves the estimate for every real time. □

Theorem PH11 (time-explicit finite-packet bound). For every real m, h, t, fixed mode n, and fixed finite packet,

$$\|U_{H_{h,n}}(t) - U_{H_n}(t)\| \le |t|\,|\delta(h,n)| \le |t|h^2 C(n),$$

$$\begin{aligned}\|f_h(t) - f(t)\|_{L^2} &\le |t| \left(\sum_{n\in S} |\delta(h,n)|^2 \|a_n\|_2^2\right)^{1/2} \\ &\le |t|h^2 \left(\sum_{n\in S} C(n)^2 \|a_n\|_2^2\right)^{1/2}.\end{aligned}$$

Consequently, for every $T \ge 0$,

$$\sup_{|t|\le T} \|f_h(t) - f(t)\|_{L^2} \le T h^2 \left(\sum_{n\in S} C(n)^2 \|a_n\|_2^2\right)^{1/2},$$

and evolution converges uniformly on every compact time interval for fixed support and initial amplitudes.

Proof. Apply the Duhamel lemma with $A = H_{h,n}$, $B = H_n$. Both are Hermitian by PH02, and PH07 gives the residual norm and bound. For each amplitude,

$$\|(U_{H_{h,n}}(t) - U_{H_n}(t))a_n\|_2 \le |t|\,|\delta(h,n)|\,\|a_n\|_2 \le |t|h^2 C(n)\|a_n\|_2.$$

Square and sum over S, apply the exact residual Parseval identity of PH08, and take the nonnegative square root. The middle coefficient is exactly the initial generator-residual norm in PH07 because these Clifford differences satisfy $\Delta_{h,n}^* \Delta_{h,n} = |\delta(h,n)|^2 I_4$. This identification is specific to the displayed symbol differences; the general Duhamel lemma is an operator-norm estimate. Replacing $|t|$ by T gives the compact-time bound. Its right side tends to zero as $h \to 0$. □

For PH10, $C(n_0) = 0$, $C(n_A) = (2\pi)^3/6$, and $C(n_D) = \sqrt{3}(2\pi)^3/6$, so

$$\left(\sum_{n\in S} C(n)^2 \|a_n\|_2^2\right)^{1/2} = \frac{(2\pi)^3}{6} = \frac{4\pi^3}{3}.$$

At its spacing, put $d = 2\pi - 6 > 0$. The axis residual has magnitude d, the diagonal residual has magnitude $\sqrt{3}d$, and their squared amplitude weights are $1/4, 1/4$. Hence the exact middle coefficient is

$$\left(\tfrac{1}{4}d^2 + \tfrac{1}{4}\,3d^2\right)^{1/2} = d,$$

and the worked bound is

$$\|f_{1/12}(t) - f(t)\|_{L^2} \le (2\pi - 6)|t| \le \frac{\pi^3}{108}|t|.$$

The triangle inequality also gives $\|f_h(t) - f(t)\|_{L^2} \le 2$ for this normalized packet, so one may use the smaller bound. Neither says that a fixed nonzero spacing has small error for arbitrarily large time. Nor does the time bound remove the cutoff: its coefficient contains high powers of mode momenta. For

arbitrary L^2 data, the Fourier-tail argument in the following part provides convergence without this uniform quadratic coefficient.

12 Probability, Hamiltonian expectation, and physical scope

Let $H_S, H_{h,S}$ denote the finite block diagonal Hamiltonians. In coefficient coordinates, squared norm and continuum Hamiltonian expectation are

$$N(a) = \sum_{n\in S} \|a_n\|_2^2, \qquad \mathcal{E}(a) = \sum_{n\in S} \langle a_n, H_n a_n\rangle,$$

with centered expectation $\mathcal{E}_h(a)$ obtained by replacing H_n with $H_{h,n}$. Hermiticity makes each expectation real: its complex conjugate is $\langle H_n a_n, a_n\rangle = \langle a_n, H_n a_n\rangle$.

Proposition PH12 (separate conservation of Hamiltonian expectation). Each exact packet evolution preserves the expectation of its own time-independent Hamiltonian:

$$\mathcal{E}(U_S(t)a) = \mathcal{E}(a), \qquad \mathcal{E}_h(U_{h,S}(t)a) = \mathcal{E}_h(a).$$

Here the packet operators act on coefficients through the isometry $\mathcal{K}_S \cong \mathcal{F}_S$. Squared-norm conservation does not identify either expectation with $N(a)$. At common initial data,

$$|\mathcal{E}_h(a) - \mathcal{E}(a)| \le h^2 \sum_{n\in S} C(n)\|a_n\|_2^2.$$

Proof. The exponential of a fixed matrix commutes with that matrix, term by term in its power series. Unitarity gives on each mode

$$\langle U_H(t)a, HU_H(t)a\rangle = \langle a, U_H(t)^* H U_H(t)a\rangle = \langle a, Ha\rangle.$$

Sum to prove both conservation laws. For common initial coefficients, Cauchy–Schwarz and PH07 give

$$|\langle a_n, (H_{h,n} - H_n)a_n\rangle| \le \|a_n\|_2 \|\Delta_{h,n} a_n\|_2 \le h^2 C(n)\|a_n\|_2^2.$$

The finite-sum triangle inequality proves the comparison. □

A unit vector in the negative-energy range of P_- has squared norm one but expectation $-E$; a unit vector in the positive-energy range has expectation E. Norm and energy are therefore different even for one mode. For e_0, the matrix gives $\langle e_0, H(m,p)e_0\rangle = m$, since the kinetic terms have zero diagonal entry in that component. PH10 has expectation 1 for both Hamiltonians. Its values $\sqrt{37}, \sqrt{109}$ are shell magnitudes, and its spinors contain both energy signs. Their expectations need not equal a positive shell value.

The analysis is free and linear. No gauge field, forcing, interaction, or time-dependent Hamiltonian occurs in these symbols. Other project modules treat different transport structures; importing them does not convert these statements into theorems about interacting fields. The centered propagator is the exact exponential of a modified spatial symbol. A numerical approximation to that exponential must be checked separately for its own norm error.

Completion and continuum evolution

13 The complete coefficient Hilbert space

Let

$$\mathcal{K} = \ell^2(\mathbb{Z}^3; \mathbb{C}^4) = \left\{ a = (a_n) : \sum_{n \in \mathbb{Z}^3} \|a_n\|^2 < \infty \right\}.$$

Its inner product is $\langle a, b \rangle = \sum_n \langle a_n, b_n \rangle$. Absolute convergence follows from the scalar Cauchy–Schwarz inequality. The norm therefore measures the total squared amplitude across all Fourier packets. A finite family of amplitudes embeds by assigning zero to every absent mode.

For a finite set $S \subset \mathbb{Z}^3$, define the truncation $P_S a$ by retaining a_n when $n \in S$ and assigning zero otherwise. Then

$$\|a - P_S a\|^2 = \sum_{n \notin S} \|a_n\|^2 .$$

The sum on the right can be made arbitrarily small because the full nonnegative series converges. Thus finite packets are dense in $\mathcal{K}$. The statement concerns the norm of the omitted amplitudes and does not depend on a particular ordering of the integer lattice.

Theorem 13.1 (Completion of a modewise isometry)**.** *Suppose that for every mode n, $V_n : \mathbb{C}^4 \to \mathbb{C}^4$ is a complex linear isometry. The coordinate rule $(Va)_n = V_n a_n$ defines a complex linear isometry on $\mathcal{K}$. If each V_n has an inverse isometry W_n, the induced V and W are inverse maps on $\mathcal{K}$.*

Proof. For each n, $\|V_n a_n\|^2 = \|a_n\|^2$. Summability of the latter terms gives membership of Va in $\mathcal{K}$, and summing the equalities gives $\|Va\| = \|a\|$. Linearity holds coordinatewise. The two inverse identities hold on each coordinate and therefore on the coefficient fields themselves. □

Apply this construction to

$$V_n(t) = e^{-itH_n}, \qquad V_{h,n}(t) = e^{-itH_{h,n}}.$$

The finite-dimensional norm preservation already proved supplies its hypothesis. Define

$$(U(t)a)_n = V_n(t)a_n, \qquad (U_h(t)a)_n = V_{h,n}(t)a_n.$$

Both maps are everywhere-defined unitary maps on $\mathcal{K}$. Their inverses are obtained at time $-t$. The modewise time-group laws imply

$$U(s+t) = U(s)U(t), \qquad U_h(s+t) = U_h(s)U_h(t).$$

This construction uses the bounded unitary map on each mode. It does not apply an unbounded Hamiltonian to an arbitrary coefficient field. In particular, the evolution is defined for more initial states than the Hamiltonian action constructed below.

13.1 The invariant finite packets

Both families commute with every finite truncation:

$$U(t)P_S = P_S U(t), \qquad U_h(t)P_S = P_S U_h(t).$$

The equality follows directly from the diagonal mode action. It says that the full evolution restricts to the finite-packet evolution already constructed. It also says that the missing tail evolves without feeding new amplitude into the retained mode set in this free model.

This property would have to be reconsidered in a model whose coefficients couple different Fourier modes. The completion argument for a diagonal multiplier makes essential use of the simple resolved action, even though the general packet language also describes nondiagonal transfers.

14 Removing the Fourier cutoff

A finite-packet consistency theorem alone says nothing about an arbitrary field with infinitely many nonzero coefficients. The missing estimate must control the error introduced by truncating that field. Exact unitarity gives a bound whose tail terms do not grow with time.

Lemma 14.1 (Two tails and a finite core)**.** *Let V and W be linear isometries on a normed space. For any a and any approximation b,*

$$\|Va - Wa\| \le 2\,\|a - b\| + \|Vb - Wb\|\,.$$

Proof. Write the residual as $V(a-b)+(Vb-Wb)+W(b-a)$. The triangle inequality bounds its norm by the sum of the three norms. Each outer term has norm $\|a-b\|$ because its operator is an isometry. □

Take $V = U_h(t)$, $W = U(t)$, and $b = P_S a$. For a fixed finite S, the compact-time packet estimate gives

$$\|U_h(t)P_S a - U(t)P_S a\| \le K_{S,a,T}h^2$$

for all $|t| \le T$ and sufficiently small positive h. The constant is allowed to depend on the support, the amplitudes, and the time window. Combining the estimates gives

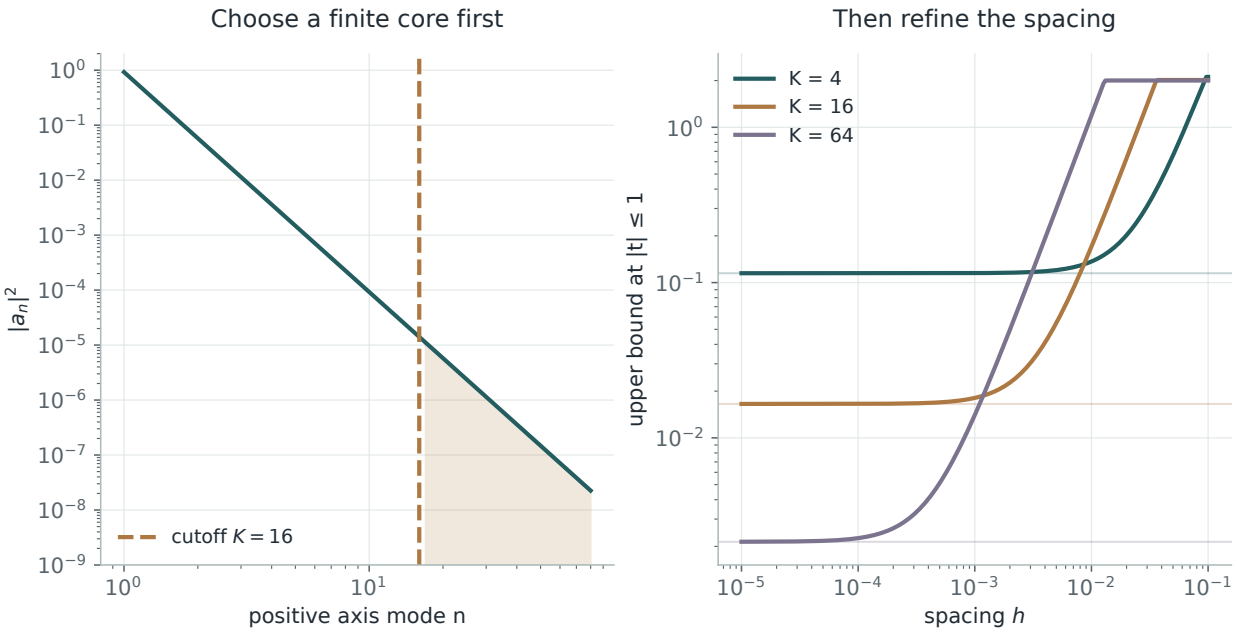
$$\|U_h(t)a - U(t)a\| \le 2\,\|a - P_S a\| + K_{S,a,T}h^2.$$

Theorem 14.2 (Compact-time strong convergence)**.** *For every $a \in \mathcal{K}$, every $T \ge 0$, and every $\varepsilon > 0$, there exists $\delta > 0$ such that*

$$0 < h < \delta, \quad |t| \le T \quad \Longrightarrow \quad \|U_h(t)a - U(t)a\| < \varepsilon.$$

Proof. Choose a finite S such that $\|a - P_S a\| < \varepsilon/4$. The finite-packet theorem gives a sufficiently small positive δ for which the core residual is less than $\varepsilon/2$ throughout the time interval. The two-tail inequality then makes the full residual less than ε. The choice of S precedes the choice of h and works for every time in the same interval. □

The quantifiers are important. One initial field and one compact time interval are fixed before the mesh scale is chosen. The theorem does not assert convergence in operator norm on the entire unit ball. It also does not assign a universal $O(h^2)$ rate to every $a \in \mathcal{K}$. For arbitrary square-summable data, a small Fourier tail may require a large cutoff, and the finite-core constant may grow with that cutoff.



The proof separates the approximation of the initial field from the consistency estimate on finitely many modes. The diagram records the order of choices in the argument.

15 Fourier synthesis on the physical torus

The scalar Fourier characters $e_n(x) = e^{2\pi i n\cdot x}$ are normalized relative to Haar measure of total mass one. For a four-spinor field, apply the scalar Fourier coefficient map to each of its four coordinates. The resulting map is

$$\mathcal{F} : \mathcal{H} \longrightarrow \mathcal{K}, \qquad \mathcal{H} = L^2(\mathbb{T}^3; \mathbb{C}^4).$$

The four-spinor Parseval identity and surjectivity of synthesis produce a complex linear isometric equivalence. In ordinary notation,

$$\|\mathcal{F} f\|_{\mathcal{K}} = \|f\|_{\mathcal{H}}, \qquad \mathcal{F}^{-1}\mathcal{F} f = f, \qquad \mathcal{F}\mathcal{F}^{-1} a = a.$$

The inverse is understood as convergence in L^2. A general field need not have a pointwise-convergent Fourier series, and no such convergence is required here.

For a finite-support coefficient field, the inverse is the familiar finite sum

$$\mathcal{F}^{-1} a = \sum_n e_n a_n.$$

The complete-space equivalence extends that reconstruction while retaining its exact norm. Its surjectivity is the step that rules out having built only a proper closed subspace of the physical field space.

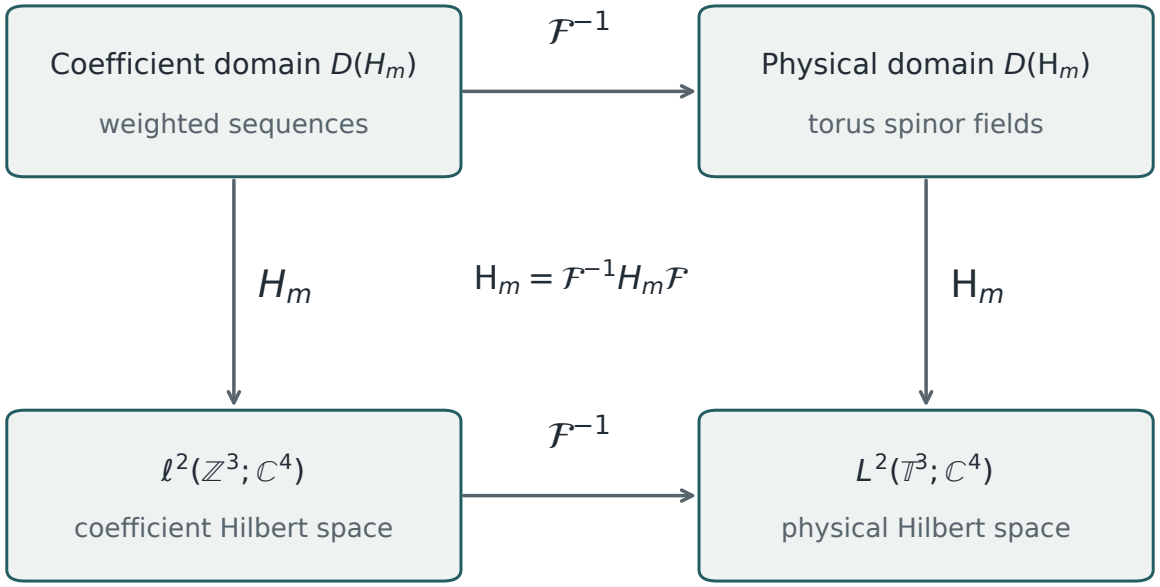
Definition 15.1 (Physical propagators)**.** Define

$$\mathcal{U}(t) = \mathcal{F}^{-1} U(t) \mathcal{F}, \qquad \mathcal{U}_h(t) = \mathcal{F}^{-1} U_h(t) \mathcal{F}.$$

Both preserve the physical L^2 norm and inner product. Their group laws follow by inserting $\mathcal{F}\mathcal{F}^{-1} = I$ between consecutive factors. Their residual satisfies the exact equality

$$\|\mathcal{U}_h(t)f - \mathcal{U}(t)f\|_{\mathcal{H}} = \|U_h(t)\mathcal{F}f - U(t)\mathcal{F}f\|_{\mathcal{K}}\,.$$

The compact-time convergence theorem therefore transfers to every $f \in \mathcal{H}$ without an additional constant.



The horizontal maps preserve inner products and norms.

Fourier analysis and synthesis change the representation while preserving the Hilbert norm. The operator in each representation is related by conjugation.

The physical torus in this theorem is a fixed flat compact space. The word physical identifies the function-space realization of the fields, rather than establishing empirical adequacy or deriving a spacetime geometry. Periodicity is an explicit boundary condition. The next part constructs the Hamiltonian on exactly this space and retains that same boundary condition.

The unbounded Hamiltonian

16 A domain defined by the operator output

The continuum mode energy grows with $|n|$. A square-summable sequence a_n need not remain square summable after multiplication by H_n. This is why the full Hamiltonian is represented as a partial linear map with a specified domain.

Define

$$\mathcal{D}(\widehat{H}_m) = \left\{ a \in \mathcal{K} : \sum_n \left\| H_n a_n \right\|^2 < \infty \right\}, \qquad (\widehat{H}_m a)_n = H_n a_n.$$

The domain is maximal for this coefficient action: it contains precisely those coefficient fields for which the stated output belongs to the ambient Hilbert space. It is a vector subspace. For example, $\|x+y\|^2 \le 2\|x\|^2 + 2\|y\|^2$ and linearity of each H_n prove closure under addition. Scalar multiplication preserves summability as well.

Every finite-support field lies in this domain, since the output again has finite support. Density of finite-support fields therefore proves density of $\mathcal{D}(\widehat{H}_m)$ in $\mathcal{K}$. Density matters because it gives uniqueness of an adjoint value when such a value exists.

The source represents this object with Mathlib's `LinearPMap`. The domain is a submodule and is carried by the operator itself. It is not replaced with a total function that assigns an arbitrary value to every vector outside the domain.

17 The adjoint domain and self-adjointness

Modewise Hermiticity gives symmetry on the maximal domain. If a, b belong to that domain, the two series in

$$\left\langle \widehat{H}_m a, b \right\rangle = \sum_n \langle H_n a_n, b_n \rangle = \sum_n \langle a_n, H_n b_n \rangle = \left\langle a, \widehat{H}_m b \right\rangle$$

converge absolutely by Cauchy–Schwarz. This establishes the inclusion $\widehat{H}_m \subseteq \widehat{H}_m^*$, including the corresponding domain inclusion.

Equality requires the converse. A merely symmetric operator can have a strictly larger adjoint domain. Here the coordinate structure makes it possible to determine that domain directly.

Theorem 17.1 (Self-adjoint maximal multiplier)**.** *The densely defined operator $\widehat{H}_m$ equals its Hilbert adjoint.*

Proof. Let $b \in \mathcal{D}(\widehat{H}_m^*)$ and write $c = \widehat{H}_m^* b$. By the adjoint relation,

$$\left\langle \widehat{H}_m a, b \right\rangle = \langle a, c \rangle \qquad (a \in \mathcal{D}(\widehat{H}_m)).$$

Fix a mode n and choose a with the single nonzero coefficient $a_n = v$, where $v \in \mathbb{C}^4$ is arbitrary. Such a field belongs to the maximal domain. The adjoint identity becomes

$$\langle H_n v, b_n \rangle = \langle v, c_n \rangle .$$

Hermiticity gives $\langle v, H_n b_n \rangle = \langle v, c_n \rangle$ for every v. Taking $v = H_n b_n - c_n$ shows $c_n = H_n b_n$. Because $c \in \mathcal{K}$, the sequence $(H_n b_n)_n$ is square summable. Thus $b \in \mathcal{D}(\widehat{H}_m)$ and the two operator actions agree there. This proves $\widehat{H}_m^* \subseteq \widehat{H}_m$. Symmetry supplies the reverse inclusion. □

The argument identifies the output before inferring membership in the maximal domain. It therefore rules out an extra adjoint-domain vector through a concrete summability requirement, rather than assuming the desired domain equality.

Closedness follows from self-adjointness. It can also be seen directly for this multiplier. If $a^{(r)} \to a$ and $\widehat{H}_m a^{(r)} \to b$ in $\mathcal{K}$, then convergence holds at each coordinate. Since the finite matrix H_n is continuous,

$$b_n = \lim_r H_n a_n^{(r)} = H_n a_n.$$

The output sequence equals $b \in \mathcal{K}$, so a belongs to the maximal domain and $\widehat{H}_m a = b$. This direct argument is included as an ordinary explanatory deduction.

18 Finite packets approximate the graph

For a domain vector a, truncation commutes with the diagonal multiplier:

$$\widehat{H}_m P_S a = P_S \widehat{H}_m a.$$

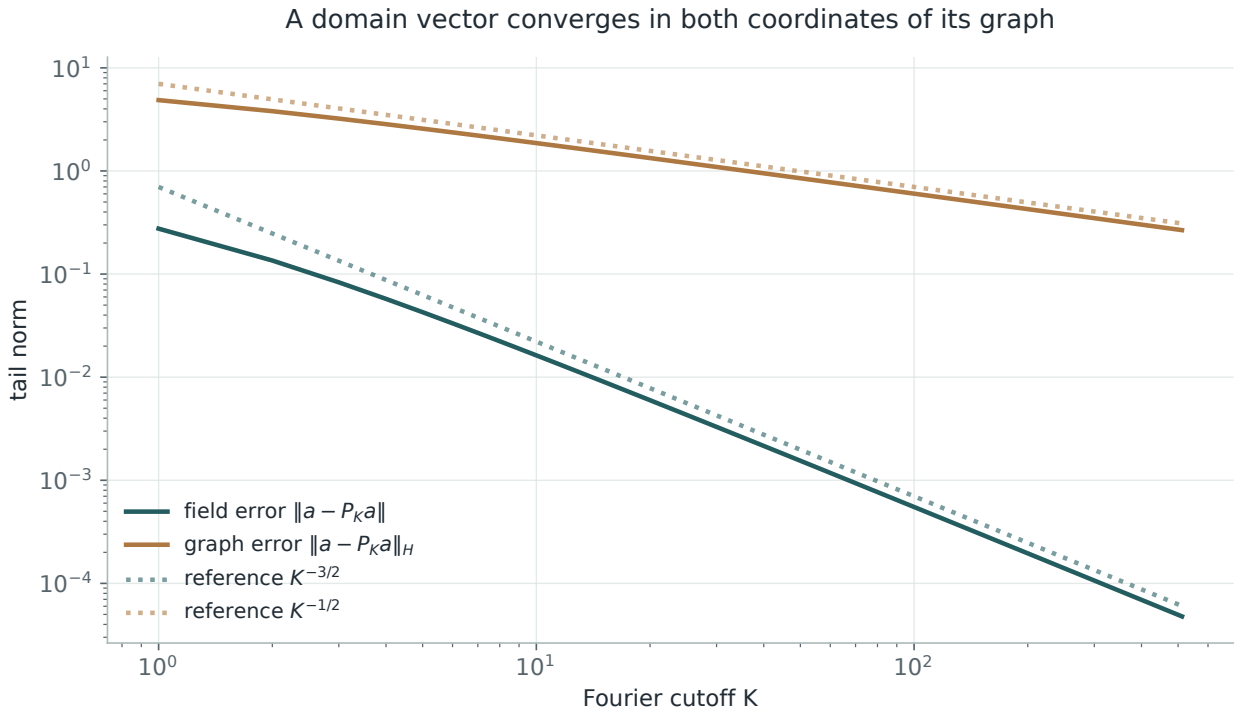
Both a and $\widehat{H}_m a$ are square summable. Given $\varepsilon > 0$, choose finite sets that make each of their tails small and take their union. The resulting one finite set S satisfies

$$\|a - P_S a\| < \varepsilon, \qquad \left\|\widehat{H}_m a - \widehat{H}_m P_S a\right\| < \varepsilon.$$

This graph-approximation statement is represented in the formal source. The finite Fourier fields are dense for the graph norm

$$\|a\|_{\text{graph}}^2 = \|a\|^2 + \left\|\widehat{H}_m a\right\|^2.$$

The term *graph core* describes this mathematical conclusion. The source proves the simultaneous approximation theorem without introducing a separate `IsCore` wrapper.



Approximating a field in its Hilbert norm and approximating it together with its Hamiltonian image are different requirements. The plotted family is a bounded illustration of the graph-tail calculation.

This result concerns truncations of one continuum operator. A lattice-to-continuum graph-convergence theorem compares different operators and still requires its own compatible-domain construction. The word graph in both phrases does not identify those two tasks.

19 The exact first-order Fourier domain

The Clifford square identity sharpens the definition of the maximal domain. Since $H_n^* = H_n$ and $H_n^2 = (m^2 + |p(n)|^2)I_4$,

$$\|H_n v\|^2 = \langle v, H_n^* H_n v\rangle = (m^2 + |p(n)|^2)\,\|v\|^2 .$$

For the coefficient field this gives the exact equivalence

$$a \in \mathcal{D}(\widehat{H}_m) \quad \Longleftrightarrow \quad \sum_n (m^2 + |p(n)|^2)\,\|a_n\|^2 < \infty.$$

The mass term is already summable because $a \in \mathcal{K}$. Membership is therefore equivalent to summability of $\sum_n |p(n)|^2\,\|a_n\|^2$.

Theorem 19.1 (Weighted domain and mass independence)**.** *For every real m,*

$$\mathcal{D}(\widehat{H}_m) = \left\{ a \in \mathcal{K} : \sum_n (1 + |p(n)|^2)\,\|a_n\|^2 < \infty \right\}.$$

Consequently $\mathcal{D}(\widehat{H}_{m_1}) = \mathcal{D}(\widehat{H}_{m_2})$ for every two real masses.

Proof. The weighted energy series is the sum of the nonnegative momentum series and $m^2 \sum_n \|a_n\|^2$. The second series is finite. Thus the energy series is finite exactly when the momentum series is finite. Adding the already finite unweighted series gives precisely the displayed inhomogeneous first-order condition. This condition contains no mass parameter. □

The inhomogeneous weight retains the zero mode even when $m = 0$. There is no loss of the ambient square-summability requirement in the massless case. The phrase first-order Fourier domain refers to this exact condition. It agrees with the conventional Fourier definition of periodic H^1 regularity; this manuscript does not claim a new checked equivalence with a separate weak-derivative implementation.

19.1 A concrete domain boundary

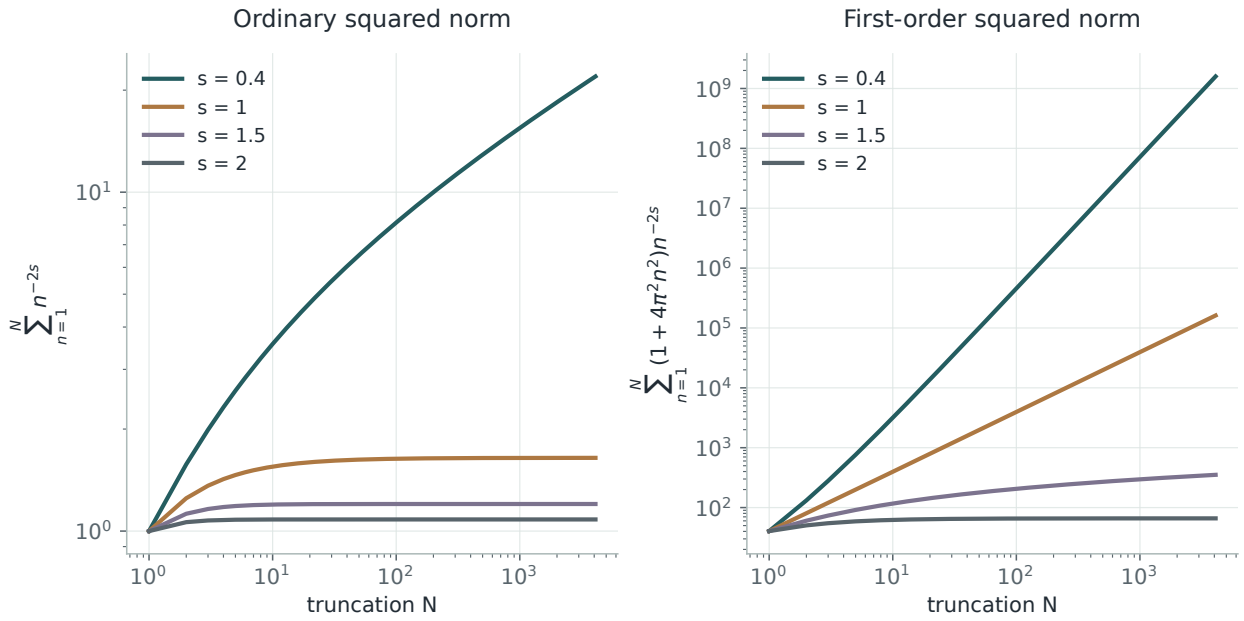
As an ordinary example, put nonzero coefficients only on the positive axis modes $n = (k, 0, 0)$, $k \geq 1$, and set $a_n = k^{-s} v$ for a fixed unit spinor v. All remaining coefficients are zero. Then

$$\sum_n \|a_n\|^2 = \sum_{k \geq 1} k^{-2s},$$

which is finite exactly when $s > 1/2$. The additional momentum series is

$$4\pi^2 \sum_{k \geq 1} k^{2-2s},$$

which is finite exactly when $s > 3/2$. Thus the choice $s = 1$ gives an L^2 field outside the Hamiltonian domain, while $s = 2$ gives one inside it. These thresholds use axis-supported data. A radially populated three-dimensional family has a different counting factor.



The norm condition and the Hamiltonian-domain condition impose different decay requirements. The exact threshold argument uses the infinite series; finite plotted partial sums only illustrate its behavior.

20 Exact operator energy

For $a \in \mathcal{D}(\widehat{H}_m)$, summing the modal norm identity gives

$$\left\|\widehat{H}_m a\right\|^2 = \sum_n (m^2 + |p(n)|^2) \left\|a_n\right\|^2 .$$

This is the squared norm of the Hamiltonian output. It is nonnegative. It must be distinguished from the expectation $\left\langle a, \widehat{H}_m a \right\rangle$, which may have either sign. The finite-packet discussion of positive and negative energy branches already exhibits that distinction.

An immediate ordinary consequence is the explicit graph norm

$$\|a\|^2_{\text{graph}} = \sum_n (1 + m^2 + |p(n)|^2) \left\|a_n\right\|^2 .$$

The graph norm is therefore equivalent to the norm defined by the first-order Fourier weight. Writing $w(n) = 1 + |p(n)|^2$, one has

$$w(n) \le 1 + m^2 + |p(n)|^2 \le (1 + m^2) w(n).$$

The domains coincide for all masses even though these graph norms and the operator actions depend on the mass.

21 The Hamiltonian on physical spinor fields

Use the checked Fourier equivalence to define

$$\mathcal{D}(H_m^{\text{phys}}) = \mathcal{F}^{-1}\mathcal{D}(\widehat{H}_m), \qquad H_m^{\text{phys}} f = \mathcal{F}^{-1}\widehat{H}_m \mathcal{F} f.$$

The second formula is used only for f in the displayed domain. Its Fourier intertwining identity is

$$\mathcal{F}(H_m^{\text{phys}} f) = \widehat{H}_m(\mathcal{F} f).$$

Theorem 21.1 (Physical torus Hamiltonian)**.** *The operator H_m^{phys} is densely defined, symmetric, self-adjoint, and closed on $\mathcal{H} = L^2(\mathbb{T}^3; \mathbb{C}^4)$. Its domain consists exactly of fields satisfying*

$$\sum_{n \in \mathbb{Z}^3} (1 + |2\pi n|^2) \left\|\hat{f}(n)\right\|^2 < \infty,$$

and for every domain vector,

$$\left\|H_m^{\text{phys}} f\right\|^2_{\mathcal{H}} = \sum_n (m^2 + |2\pi n|^2) \left\|\hat{f}(n)\right\|^2 .$$

Proof. The domain characterization is the coefficient characterization applied to $\mathcal{F} f$. Density transports through the isometric equivalence. For two domain vectors, transporting the inner product and using coefficient-space symmetry proves physical-space symmetry.

To prove the reverse adjoint-domain inclusion, let g have physical adjoint value z. For every coefficient domain vector a, test the physical adjoint identity at $\mathcal{F}^{-1}a$. Fourier isometry gives

$$\left\langle \widehat{H}_m a, \mathcal{F}\, g \right\rangle = \langle a, \mathcal{F}\, z \rangle \,.$$

Thus $\mathcal{F}\, g$ belongs to the coefficient adjoint domain with adjoint value $\mathcal{F}\, z$. Coefficient self-adjointness puts $\mathcal{F}\, g$ in the maximal coefficient domain, hence puts g in the physical domain. The actions agree. Symmetry supplies the opposite inclusion, giving equality with the adjoint. Closedness follows, and Fourier isometry transports the exact energy sum. □

On finite Fourier polynomials, the operator agrees with the matrix differential expression $-i \sum_j \alpha_j \partial_j + m\beta$ already calculated. Its extension is specified by Fourier transport and the maximal domain, so no pointwise differentiability assumption is imposed on every L^2 representative. Equality with a separately constructed weak differential operator is a further interface statement if that alternate realization is introduced.

The operator and its evolution

22 The completed operator construction

The free field now has two precise mathematical descriptions. Modewise exponentials define a unitary evolution on the complete coefficient space and, by Fourier conjugation, on physical torus fields. The maximal coefficient multiplier defines a self-adjoint unbounded Hamiltonian whose domain is the first-order weighted Fourier space.

These constructions answer different questions. Norm preservation determines the behavior of the bounded evolution maps. Self-adjointness determines the operator and its adjoint domain. Finite-packet graph approximation controls the state and its Hamiltonian image at the same time. No one of these statements can replace the others.

Construction	Mathematical role
Clifford mode matrices	Exact mass shell and finite-dimensional self-adjointness.
Mode exponentials	Group identities and unitary spinor dynamics.
Completed coefficient evolution	Removal of finite support by norm stability.
Fourier equivalence	Lossless transport between coefficients and physical fields.
Maximal Hamiltonian	Density, adjoint-domain equality, and closedness.
Weighted domain	A mass-independent characterization of first-order regularity.

The explicit spectral-projector formula, sharp finite Duhamel coefficient, and worked examples are deductions with proofs in the text. The technical appendix distinguishes these deductions from their exact formal correspondences. Recorded compiler evidence applies to the named statements and hypotheses, not to every sentence of interpretation.

23 Connecting the group to its generator

The next chapter proves that the constructed group is the spectral evolution of this Hamiltonian:

$$\mathcal{U}(t) = e^{-itH_m^{\mathrm{phys}}}.$$

The proof establishes strong continuity and differentiates on the full Hamiltonian domain with control of the Fourier tail. The reverse domain inclusion then shows that the infinitesimal generator has exactly that domain. The equality is therefore a conclusion about the complete Hilbert-space group, not an inference from agreement on finitely many modes.

The full-domain generator identification has a checked Lean implementation. The abstract spectral-functional-calculus identity remains a written deduction without a separate formal API linkage. The technical appendix distinguishes their verification coverage; no additional analytic hypothesis is introduced.

24 Operator limits and regularity

The centered family acts on the same infinite Fourier index set as the continuum family. The generator chapter proves convergence on the Hamiltonian domain and strong-resolvent convergence. It also gives a quadratic spacing estimate for the propagator, uniformly on bounded time intervals, under an explicit third-order spatial Fourier moment condition.

The regularity condition is essential to that estimate. Arbitrary square-integrable data retain the strong compact-time convergence statement, but are not assigned a universal quadratic rate. The first-order Hamiltonian domain alone is not substituted for the third-order moment used by the quantitative bound.

A changing finite spatial grid is a different comparison problem. It needs sampling, aliasing control, embeddings into the physical Hilbert space, and compatibility with geometric refinement. The centered multiplier family does not supply those ingredients merely by using a spacing parameter.

25 Beyond the free field

A prescribed background connection changes the operator and its domain. Dynamical gauge fields additionally require field energy, a Gauss constraint, matter backreaction, and an existence theory. The gauge chapter develops the finite transport and refinement ingredients for these questions without treating them as a coupled continuum evolution theorem.

The gravity, arithmetic refinement, and cosmology chapters retain their own mathematical and physical assumptions. The free model provides a controlled analytic sector to which such constructions may be compared. It neither derives a dynamical metric nor predicts an observational parameter.

Sources and further reading

[1] Komorebi P., *Formal source for the free Dirac construction*. Exact module, compiler, and archive identities are recorded in the formal verification appendix. Its certificate applies to the stated source declarations and hypotheses.

[2] Komorebi P., *Fourier Packets and Hilbert Stability for Dirac Evolution*. The finite-packet results PH01 to PH12 provide the underlying notation. Their written deductions and exact formal correspondences are distinguished in the verification appendix.

[3] Komorebi P., *Projector Refinement Physics*: foundational and cosmological studies, technical supplement, selected second implementation, and covariant backreaction analysis. The technical appendices identify their source roles and computational evidence.

[4] P. A. M. Dirac, *The quantum theory of the electron*, Proceedings of the Royal Society of London. Series A **117** (1928), no. 778, 610–624. doi:10.1098/rspa.1928.0023.

[5] C. Roiesnel, *Definition of the covariant lattice Dirac operator*, Physical Review D **87** (2013), 074505. doi:10.1103/PhysRevD.87.074505. The symmetric difference stencil is also recorded in Eq. (9) of the author's preprint arXiv:1211.1944v1 (8 November 2012).

[6] P. Bader, S. Blanes, F. Casas, and M. Seydaoğlu, *An efficient algorithm to compute the exponential of skew-Hermitian matrices for the time integration of the Schrödinger equation*, Mathematics and Computers in Simulation **194** (2022), 383–400. doi:10.1016/j.matcom.2021.12.002. Author preprint: arXiv:2103.10132v2 (7 December 2021).

The Dirac Group and Its Generator

Spectral identification, strong resolvents, and quantitative centered approximation

The preceding chapter constructed a self-adjoint free Dirac Hamiltonian and a modewise unitary evolution. This chapter proves that they describe the same dynamics. The infinitesimal generator is exactly $-iH$, and its domain is precisely the maximal Hamiltonian domain.

The centered Fourier generators also converge on that domain and in the strong resolvent sense. A finite third-order spatial Fourier moment supplies a quadratic spacing bound for the propagator, uniform on each bounded time interval. The argument concerns a fixed flat torus and an unchanged infinite Fourier index set.

The strong generator and centered graph, resolvent and third-moment estimates have checked Lean implementations with physical-space transport. The abstract spectral-functional-calculus identity remains a written deduction without a formal API linkage. These results do not identify the model with a physical particle or an interacting continuum theory.

1 The Hilbert space and the Hamiltonian

Let $\mathfrak{h} = \mathbb{C}^4$ with its standard Hermitian norm, and set

$$\mathcal{H} = \ell^2(\mathbb{Z}^3; \mathfrak{h}), \qquad k_n = 2\pi n, \qquad n \in \mathbb{Z}^3. \tag{1}$$

Fix the Hermitian Dirac matrices $\beta, \alpha_1, \alpha_2, \alpha_3$ with

$$\beta^2 = \mathrm{I}, \quad \alpha_j^2 = \mathrm{I}, \quad \alpha_j\alpha_l + \alpha_l\alpha_j = 0 \ (j \neq l), \quad \alpha_j\beta + \beta\alpha_j = 0. \tag{2}$$

All masses $m \in \mathbb{R}$ are permitted. Define the four by four mode matrix

$$H_m(n) = m\beta + \sum_{j=1}^{3} k_{n,j}\alpha_j, \qquad H_m(n)^2 = (m^2 + |k_n|^2)\mathrm{I}. \tag{3}$$

The last equality is an exact Clifford identity. Let $H = H_m$ denote the maximal diagonal multiplier on $\mathcal{H}$:

$$\begin{aligned} \mathcal{D}(H) &= \left\{ f \in \mathcal{H} : \sum_n |k_n|^2 \, \|f_n\|^2 < \infty \right\}, \\ (Hf)_n &= H_m(n) f_n. \end{aligned} \tag{4}$$

The preceding operator construction establishes self adjointness, closedness, and density, and identifies $\mathcal{D}(H)$ with the weighted first order Fourier domain. The equivalent inhomogeneous definition uses $1 + |k_n|^2$. The domain does not depend on m.

Define the unitary family previously constructed mode by mode:

$$(U(t)f)_n = \exp(-itH_m(n))f_n. \tag{5}$$

As in the earlier modules, the complex inner product is linear in its second coordinate. Equations (3) and (5) fix the sign of the generator to be $-iH$.

1.1 Why the distinction matters

An exact matrix identity on each mode need not automatically justify an infinite dimensional derivative, because the norms $\|H_m(n)\|$ grow with $|n|$. Nor does equality on finite Fourier packets alone identify the maximal generator. The next proof controls the tails through the natural graph norm. This supplies the missing step between the checked propagator and the unbounded Hamiltonian.

2 The continuum group and its maximal generator

Lemma 2.1 (Strong continuity)**.** *The family $(U(t))_{t\in\mathbb{R}}$ is a strongly continuous unitary group on $\mathcal{H}$. Its action commutes with every finite Fourier truncation.*

Proof. Every $H_m(n)$ is Hermitian, so its exponential is unitary. The group law and $U(0) = \mathrm{I}$ hold at each n and hence in $\mathcal{H}$. For any finite $S \subset \mathbb{Z}^3$, let P_S be coefficient truncation. Diagonality gives $U(t)P_S = P_S U(t)$ exactly. For $f \in \mathcal{H}$, select S so that $\|f - P_S f\| < \varepsilon/4$. At any t,

$$\|U(t)f - f\| \le 2\,\|f - P_S f\| + \|U(t)P_S f - P_S f\|\,.$$

The second term tends to zero as $t \to 0$, being the norm of a finite tuple of continuous matrix exponentials. The first is smaller than $\varepsilon/2$. Thus $U(t)f \to f$ for every f. Continuity at other times follows from the group law and unitarity. □

Lemma 2.2 (Domain preservation and graph isometry)**.** *For every $t \in \mathbb{R}$, $U(t)\mathcal{D}(H) = \mathcal{D}(H)$, and on this domain*

$$HU(t)f = U(t)Hf, \qquad \|U(t)f\|^2 + \|HU(t)f\|^2 = \|f\|^2 + \|Hf\|^2\,.$$

Proof. The exponential of a fixed matrix commutes with that matrix. Consequently $H_m(n)U_n(t)f_n = U_n(t)H_m(n)f_n$. Since $U_n(t)$ preserves the spinor norm, $\sum_n \|H_m(n)U_n(t)f_n\|^2 = \sum_n \|H_m(n)f_n\|^2$. The series is finite precisely on $\mathcal{D}(H)$. Replacing t by $-t$ gives equality of domains, while equality of the two graph norms follows from modewise unitarity. □

Theorem 2.3 (Maximal strong generator)**.** *The strong generator A of $U(t)$ has domain exactly $\mathcal{D}(H)$ and acts as $\mathsf{A} = -iH$:*

$$\mathcal{D}(\mathsf{A}) = \mathcal{D}(H), \quad \lim_{t\to 0}\left\|\frac{U(t)f - f}{t} + iHf\right\| = 0 \quad (f \in \mathcal{D}(H)). \tag{6}$$

Conversely, if the strong difference quotient has a limit for $f \in \mathcal{H}$, then $f \in \mathcal{D}(H)$ and that limit equals $-iHf$.

Proof. Fix $f \in \mathcal{D}(H)$. For each n, the ordinary finite dimensional exponential identity gives

$$\frac{U_n(t) - \mathrm{I}}{t} f_n = -i\int_0^1 U_n(st)H_m(n)f_n\,ds \quad (t \neq 0). \tag{7}$$

The integral norm is bounded by $\|H_m(n)f_n\|$ using unitarity. It converges to $-iH_m(n)f_n$ as $t \to 0$. More precisely, after adding $iH_m(n)f_n$, the residual is bounded by $2\,\|H_m(n)f_n\|$ for every $t \neq 0$. Since $\sum_n \|H_m(n)f_n\|^2$ is finite, dominated convergence for nonnegative series implies that the squared sum of the residual norms tends to zero. This is exactly the asserted strong derivative.

For the converse suppose $(U(t)f - f)/t \to g$ in $\mathcal{H}$. Evaluation at a fixed Fourier index n is a continuous linear map from $\mathcal{H}$ to $\mathfrak{h}$, and hence the corresponding finite dimensional difference quotient converges to g_n. Its unique limit is $-iH_m(n)f_n$. Thus $g_n = -iH_m(n)f_n$ for every n. Since $g \in \mathcal{H}$, the sequence $(H_m(n)f_n)$ is square summable and $f \in \mathcal{D}(H)$. The operator and its entire domain have now been identified, not only their restrictions to a core. □

Corollary 2.4 (Spectral exponential and physical evolution)**.** *The checked unitary group equals the spectral exponential of the maximal self adjoint multiplier:*

$$U(t) = e^{-itH} \qquad (t \in \mathbb{R}).$$

Let $\mathcal{F} : L^2(\mathbb{T}^3;\mathbb{C}^4) \to \mathcal{H}$ be the existing unitary Fourier equivalence, and let $H^{\mathrm{phys}} = \mathcal{F}^{-1}H\mathcal{F}$ with transported domain. Then the previously constructed physical torus group obeys

$$U^{\mathrm{phys}}(t) = \mathcal{F}^{-1}U(t)\mathcal{F} = e^{-itH^{\mathrm{phys}}}.$$

Proof. By Lemma 2.1 the group is strongly continuous and unitary. By Theorem 2.3, its full strong generator is $-iH$. The uniqueness assertion of the Stone theorem, or equivalently the spectral theorem for the self adjoint H coupled with uniqueness of strongly continuous unitary evolutions, identifies it with e^{-itH}. Functional calculus is preserved under unitary conjugation. Applying $\mathcal{F}^{-1}(\cdot)\mathcal{F}$ proves the physical identity together with the transported domain statement. □

3 Centered generator family and resolvent limit

For each nonzero real h, let

$$q_h(x) = \frac{\sin(hx)}{h}, \qquad H_{m,h}(n) = m\beta + \sum_{j=1}^{3} q_h(k_{n,j})\alpha_j. \tag{8}$$

Write H_h for the modewise centered operator on the *same* coefficient space $\mathcal{H}$. Its exponential $U_h(t)$ is the earlier centered coefficient propagator. There is no projection onto a changing sampled lattice in this definition.

Lemma 3.1 (Bounded centered operator)**.** *For fixed $h \neq 0$, H_h is a bounded self adjoint operator on all of $\mathcal{H}$, and*

$$\|H_h\| \leq \sqrt{m^2 + 3h^{-2}}. \tag{9}$$

Proof. Each mode matrix is Hermitian and its square is $(m^2 + \sum_j q_h(k_{n,j})^2)\mathrm{I}$ by the same Clifford relations. Since $|q_h(x)| \leq 1/|h|$, its operator norm is bounded by the right side of (9), uniformly over all modes. A bounded diagonal operator with Hermitian mode matrices is self adjoint on $\mathcal{H}$. □

Theorem 3.2 (Full graph-domain convergence)**.** *For every $f \in \mathcal{D}(H)$, the centered generators converge in the norm of $\mathcal{H}$:*

$$\lim_{h \to 0,\, h \neq 0} \|H_h f - Hf\| = 0. \tag{10}$$

No regularity beyond the checked first order Fourier domain is needed.

Proof. Write $\delta_{h,j}(n) = q_h(k_{n,j}) - k_{n,j}$. By the Clifford identities, the self adjoint difference matrix $\Delta_h(n) = H_{m,h}(n) - H_m(n) = \sum_j \delta_{h,j}(n)\alpha_j$ satisfies

$$\|\Delta_h(n)\psi\|^2 = \sum_{j=1}^{3} |\delta_{h,j}(n)|^2 \, \|\psi\|^2 . \tag{11}$$

For every fixed n, $\delta_{h,j}(n) \to 0$. The elementary bound $|\sin x| \le |x|$ gives $|q_h(k)| \le |k|$ and hence $|\delta_{h,j}(n)| \le 2|k_{n,j}|$. Equation (11) therefore bounds each squared modal difference by $4|k_n|^2 \, \|f_n\|^2$, summable by (4). Dominated convergence of this series yields (10). □

Theorem 3.3 (Strong resolvent convergence). *For each $z \in \mathbb{C} \setminus \mathbb{R}$ and every $f \in \mathcal{H}$,*

$$(H_h - z\mathrm{I})^{-1} f \longrightarrow (H - z\mathrm{I})^{-1} f \quad \text{in } \mathcal{H}, \qquad h \to 0. \tag{12}$$

Proof. Both H_h and H are self adjoint, so their resolvents exist off the real axis and satisfy $\|(H_h - z\mathrm{I})^{-1}\| \le |\operatorname{Im} z|^{-1}$. Set $g = (H - z\mathrm{I})^{-1} f$. By the defining property of the unbounded resolvent, $g \in \mathcal{D}(H)$. As H_h is bounded everywhere, the resolvent identity applied to g is legitimate:

$$(H_h - z\mathrm{I})^{-1} f - (H - z\mathrm{I})^{-1} f = (H_h - z\mathrm{I})^{-1}(H - H_h)g.$$

Its norm is bounded by $|\operatorname{Im} z|^{-1} \, \|(H - H_h)g\|$, which tends to zero by Theorem 3.2. This establishes the strong resolvent limit for every f, not just finite packets. □

4 A global quadratic spacing estimate

Define a third order spatial Fourier moment whenever the series exists:

$$M_3(f)^2 = \sum_{n \in \mathbb{Z}^3} \left(\sum_{j=1}^{3} |k_{n,j}|^6 \right) \|f_n\|^2 . \tag{13}$$

It is controlled by the usual inhomogeneous H^3 Fourier norm. Its exact form is useful because it matches the individual centered derivative errors.

Lemma 4.1 (Uniform sine error). *For every real x and every $h \neq 0$,*

$$\left| \frac{\sin(hx)}{h} - x \right| \le \frac{h^2}{6} |x|^3. \tag{14}$$

Proof. Taylor's theorem with integral remainder gives $|\sin y - y| \le |y|^3/6$ for every real y. Indeed the second derivative of $\sin y$ is $-\sin y$, with $|\sin y| \le |y|$, so the integrated remainder is at most $|y|^3/6$. Substitute $y = hx$ and divide by $|h|$. □

Theorem 4.2 (Generator and propagator rates). *If $M_3(f) < \infty$, then for every real $h \neq 0$,*

$$\|(H_h - H)f\| \le \frac{h^2}{6} M_3(f), \tag{15}$$

and for every $t \in \mathbb{R}$,

$$\left\| \left(U_h(t) - U(t)\right) f \right\| \leq \frac{|t| h^2}{6} M_3(f). \tag{16}$$

In particular, for every $T \geq 0$, the error is at most $Th^2 M_3(f)/6$ uniformly on $|t| \leq T$.

Proof. Substituting the pointwise estimate (14) into the exact Clifford equality (11), summing over n, and taking square roots proves (15).

At a fixed mode n, the finite dimensional Duhamel identity gives, for $t \geq 0$,

$$U_{h,n}(t) - U_n(t) = -i \int_0^t U_{h,n}(t-s) \left(H_{m,h}(n) - H_m(n)\right) U_n(s)\, ds.$$

The same norm estimate holds for negative t by reversing integration limits. Both surrounding propagators are unitary. From (11), the operator norm of the middle factor equals $\left(\sum_j |\delta_{h,j}(n)|^2\right)^{1/2}$. Therefore

$$\left\| \left(U_{h,n}(t) - U_n(t)\right) f_n \right\| \leq |t| \left(\sum_j |\delta_{h,j}(n)|^2 \right)^{1/2} \|f_n\|.$$

The squared bound is summable when $M_3(f)$ is finite. Summing and applying (14) proves (16). The argument does not assume that the two different Dirac matrices commute, a property they generally lack. □

Corollary 4.3 (Physical torus norm estimates). *Let $\phi \in L^2(\mathbb{T}^3; \mathbb{C}^4)$ and $f = \mathcal{F}\phi$. Every convergence and estimate in Theorems 3.2, 3.3, and 4.2 holds after unitary Fourier transport. In particular, if $M_3(\mathcal{F}\phi) < \infty$ and $T \geq 0$, then*

$$\sup_{|t| \leq T} \left\| U_h^{\mathrm{phys}}(t)\phi - U^{\mathrm{phys}}(t)\phi \right\| \leq \frac{Th^2}{6} M_3(\mathcal{F}\phi).$$

Proof. The torus Fourier map preserves Hilbert norms and carries the centered and continuum operator actions to the respective coefficient multipliers. Apply the preceding results and then take norms after conjugation. For the resolvent, unitary conjugation preserves inverses on their transported domains. □

5 A sharper distinction between three limits

The mathematical conclusions concern three different topologies and two different regularity classes. Their separation is essential to the meaning of the estimates and to any formal implementation.

Statement	Input	Conclusion
Generator identification	$f \in \mathcal{D}(H)$	Strong derivative at 0
Centered graph limit	$f \in \mathcal{D}(H)$	$H_h f \to Hf$
Strong resolvents	All $f \in \mathcal{H}$	Resolvent action in norm
Compact time rate	$M_3(f) < \infty$	$O(h^2)$ on $\|t\| \leq T$
General evolution	All $f \in \mathcal{H}$	Strong convergence, no global rate

The bounded centered operators on the fixed Fourier index space are *not* the same object as a physical finite grid generator. A full spatial lattice theorem would need sampling and interpolation maps, a Nyquist projection, an aliasing estimate, and compatibility with the geometry and gauge refinement infrastructure. Neither Theorem 3.3 nor 4.2 supplies those identifications without further work.

6 Worked model and reproducible numerical checks

For $n = (1, 0, 0)$ and mass $m = 1$,

$$H(n) = \beta + 2\pi\alpha_1, \qquad H(n)^2 = (1 + 4\pi^2)\mathrm{I},$$

so with $E = \sqrt{1 + 4\pi^2}$ the continuum propagator has the exact closed form

$$e^{-itH(n)} = \cos(Et)\mathrm{I} - i\frac{\sin(Et)}{E}H(n).$$

For the centered model, replace 2π by $\sin(2\pi h)/h$. Both formulas hold by grouping the even and odd powers of the matrix exponential, with the removable $E = 0$ value handled continuously. This gives a direct sign check at $t = 0$.

The accompanying Python regression is a finite model with $15^3 = 3375$ Fourier modes and $m = 0.82$. Its complex test field decays with the power $(1 + |k|^2)^{-2.83}$, and it uses exact Clifford matrix identities rather than a generic time-step solver. Results include L^2 generator errors, propagator errors, resolvent errors, difference quotients, and analytic upper bounds. The figures are produced by the supplied source script; they are not evidence for the infinite dimensional theorems.

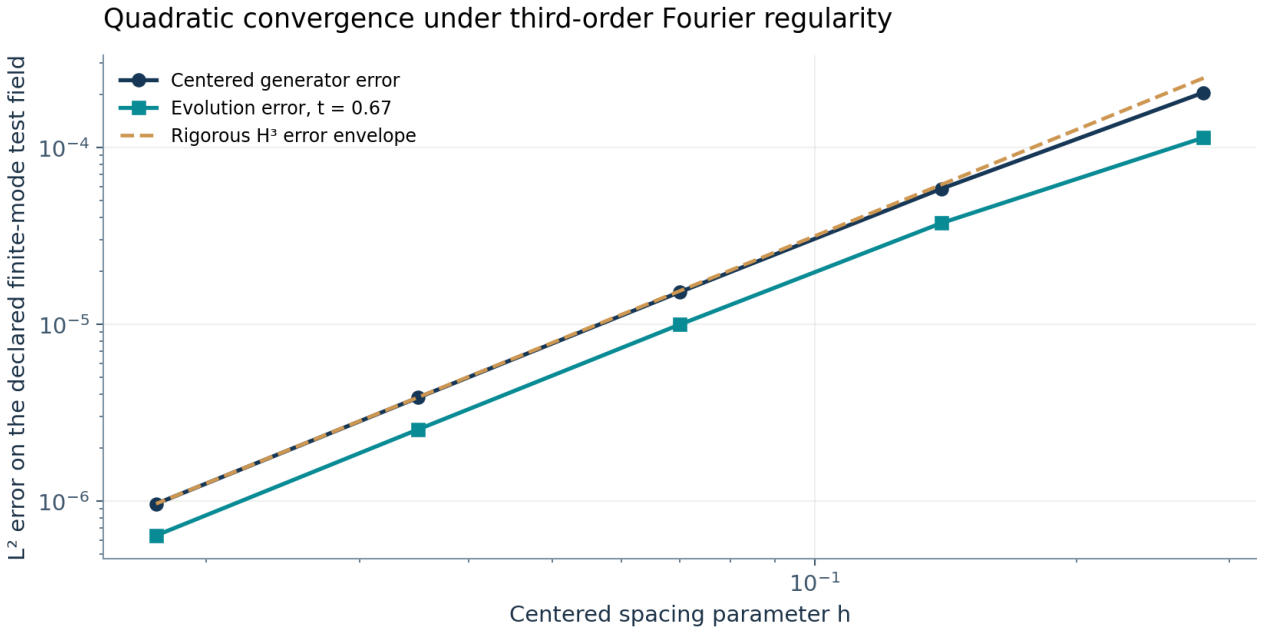


Figure 1. Finite Fourier test of the centered generator and propagator. The dashed curve is the $h^2 M_3/6$ analytic envelope for the generator, rather than a fitted convergence law.

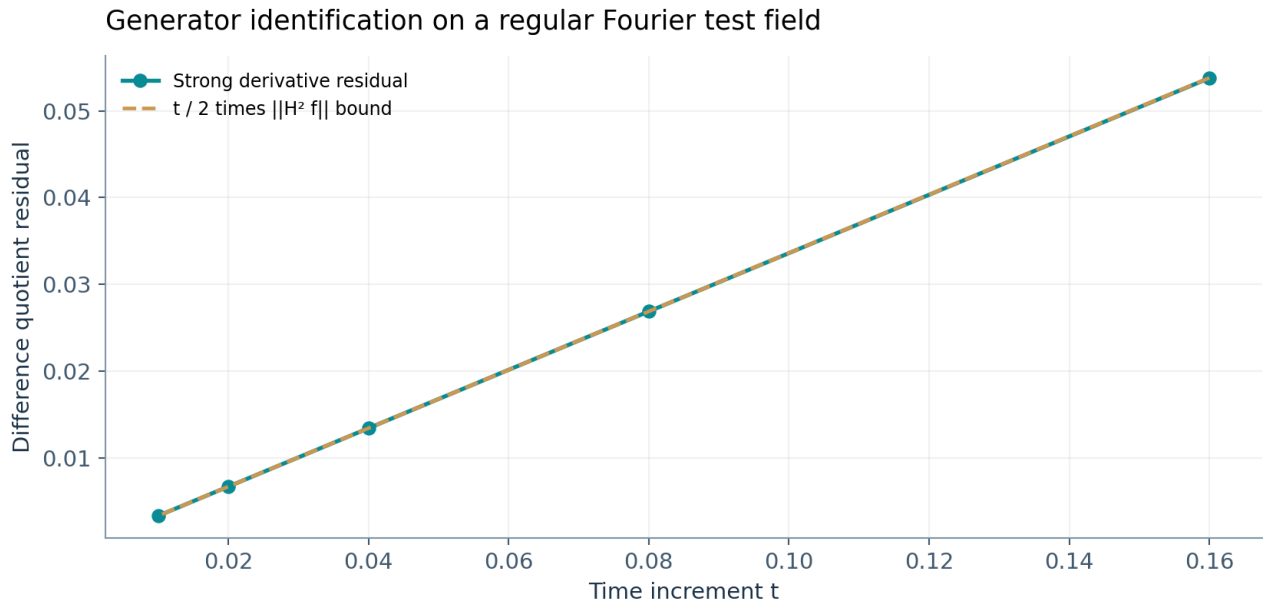


Figure 2. Difference quotient against $-iHf$ on the regular test field. The displayed bound is $|t| \left\|H^2 f\right\| /2$, the ordinary finite mode Taylor estimate.

The script additionally rejects a deliberate negative control: replacing e^{-itH} by e^{+itH} while keeping the proposed generator $-iH$ produces a difference quotient residual that approaches $2\,\|Hf\|$ rather than zero. This guards the sign convention in the construction.

7 Scope and connection to gauge transport

The generator and resolvent arguments complete the analytic free-field construction on the fixed torus coefficient space. The strong generator, its complete domain, the physical Schrödinger equation, centered graph and resolvent limits, and third-order moment bounds have compiler-checked counterparts. The abstract spectral-functional-calculus API linkage remains separate. The technical appendix gives the exact correspondence and verification boundary.

A centered Fourier symbol on the unchanged infinite mode set is not a changing finite spatial lattice. A lattice comparison must supply sampling and interpolation, control aliasing, and prove compatibility with the chosen geometry. A background gauge coupling must also specify its operator domain and evolution. Dynamical gauge and gravitational fields impose additional equations and constraints.

The following gauge treatment begins with those local transport structures. It does not turn the free spectral result into an interacting existence theorem. The production and backreaction obstruction in the matter model is likewise unaffected by this free-field analysis.

PART II

Gauge transport and gravity

A compatible link law is not yet a coupled field evolution. Gauge transport introduces ordered paths and curvature. Gravity adds dynamical geometry, constraints, boundary conditions, and a distinct continuum comparison.

IN THIS PART

Gauge Covariance and Geometric Refinement 75

Constrained Gravity and Strong Fields 107

Gauge Covariance and Geometric Refinement

Ordered transport, curvature, and the control of changing resolution

Transport beyond the free field

The free Dirac construction resolves a translation-invariant field into independent Fourier modes. Gauge coupling changes that organization. Transport between neighboring fibers carries additional information, and a comparison between resolutions must preserve the order and orientation of the links.

A link has a named source and target. Its transformation law is fixed by that direction. Composition cancels the gauge factors at internal endpoints, while a closed path retains conjugation at its basepoint. This elementary mechanism supports the treatment of holonomy, curvature, covariant differences, and refinement.

The finite Maxwell–Dirac example combines matter continuity, electric work, and propagation of the Gauss constraint. Geometric interpolation then separates edge compatibility, curvature defects, and Hilbert normalization. Regularity assumptions turn these structural conditions into quantitative estimates.

The constructions concern specified finite systems and prescribed background connections. A coupled continuum gauge theory additionally requires an operator domain, existence of the field evolution, propagation of constraints, and a convergence theorem in an explicitly chosen topology. Local covariance does not discharge those obligations.

The visual atlas develops the finite examples. The technical appendix records formal source coverage and the limits of its verification evidence. The argument here remains readable without following implementation identifiers.

Reading routes and conventions

Begin with endpoint covariance if gauge theory is new to you. The cancellation of intermediate frame changes explains both open-path covariance and closed-loop conjugation. In an Abelian group, conjugation disappears; in a non-Abelian group, the trace or conjugacy class is the natural loop observable. The ordered product always records which link acts first.

For finite dynamics, continue to the semidiscrete matter–electric-field system. Here continuity, preservation of a Gauss residual, and energy exchange follow from a specified finite system of equations. The distinction between a time-dependent frame change and the residual symmetry after choosing temporal gauge is made explicit. The consistency calculation also states how a smooth connection is converted into a link; its sampling rule affects the order of accuracy.

For refinement, follow the second part from actual half-edge geometry to cubical paths and Hilbert norms. The text keeps the straight-edge mismatch separate from the curvature contribution. Regular line-integral links remove the former and bound the latter. The reconstruction discussion then explains what is known about state spaces and which generator or evolution comparisons still have to be proved.

Convention	Meaning
Inner product	$\langle v, w\rangle$ is conjugate-linear in v and linear in w. Matrix adjoints are written T^*.
Endpoint transport	$T_{y\leftarrow x}$ maps the fiber at x to the fiber at y. Under a local unitary frame g, it becomes $g_y T_{y\leftarrow x} g_x^{-1}$.
Matrix order	$T_{z\leftarrow y}T_{y\leftarrow x}$ transports first from x to y, then from y to z. The rightmost factor acts first.
Linked shifts	A coefficient U_j in $U_j\psi_{j+1}$ transports from $j+1$ back to j. Thus $U'_j = g_j U_j g_{j+1}^{-1}$. An outward link uses the opposite endpoint convention.
Gauge groups	$U(1)$ links are unit complex numbers. The finite non-Abelian construction uses unitary matrices, with $SU(2)$ as an explicit example. Scalar multiplication identities are not silently generalized to noncommuting matrices.
Mesh and time	Mesh widths are positive when dividing by them. The finite matter system uses $\hbar = 1$ and the displayed charge convention. Refinement level and evolution time are different parameters.
Norm and measure	Spinor and operator norms are specified in the statement. Half-edge and cubical interpolation use their stated cell weights; an unweighted duplication map is not automatically isometric.
Background field	A prescribed connection supplies link integrals and curvature estimates. A dynamical gauge field additionally needs its own evolution and constraints.

The geometric source modules also use forward-link notation. The source notes translate it explicitly into the endpoint convention above. In particular, the word “first” in an edge-subdivision structure is read together with its product equation and endpoint convention.

Gauge transport and finite dynamics

1 Links compare neighboring fibers

A field assigns a vector to each site, but a local change of frame can be chosen independently at every site. A difference such as $\psi_{x+\widehat{\mu}} - \psi_x$ therefore needs a rule for comparing its two entries. A link supplies that rule. Its orientation is part of its definition, rather than a sign convention to be inferred after a calculation.

Let V be a finite periodic lattice and let every fiber be a copy of $\mathbb{C}^r$. We initially allow any unitary gauge group acting on this fiber. In the convention used for the spatial stencils below,

$$U_{x,\mu} : \mathcal{H}_{x+\widehat{\mu}} \longrightarrow \mathcal{H}_x, \qquad U_{x,\mu}^* U_{x,\mu} = U_{x,\mu} U_{x,\mu}^* = I. \tag{1}$$

The gauge $g = (g_x)_{x\in V}$ changes coordinates by $\psi_x^g = g_x \psi_x$. Requiring the transported vector to have the frame of its destination gives

$$U_{x,\mu}^g = g_x U_{x,\mu} g_{x+\widehat{\mu}}^{-1}, \qquad U_{x,\mu}^g \psi_{x+\widehat{\mu}}^g = g_x U_{x,\mu} \psi_{x+\widehat{\mu}}. \tag{2}$$

Products act on column vectors from right to left. More generally, a transporter $T_{y\leftarrow x}$ has the transformation $g_y T_{y\leftarrow x} g_x^{-1}$. This notation checks both the input and output of every product.

The forward link is

$$u_{x,\mu} = U_{x,\mu}^{-1} : \mathcal{H}_x \longrightarrow \mathcal{H}_{x+\widehat{\mu}}, \qquad u_{x,\mu}^g = g_{x+\widehat{\mu}} u_{x,\mu} g_x^{-1}. \tag{3}$$

The older time-dependent Abelian walk writes this forward link as L_j. The later geometric Lean sources use u. The non-Abelian stencil and the semidiscrete Hamiltonian use the backward U. These formulas describe the same edge in opposite directions. In particular, reversing an edge changes its transporter to the inverse, not to the same matrix in a different notation.

Proposition 1.1 (Covariance of linked differences)**.** *For spacing $h > 0$, define*

$$(D_{\mu,U}^+ \psi)_x = \frac{U_{x,\mu}\psi_{x+\widehat{\mu}} - \psi_x}{h}, \tag{4}$$

$$(\Delta_U \psi)_x = \frac{1}{h^2} \sum_\mu \left(U_{x,\mu}\psi_{x+\widehat{\mu}} + U_{x-\widehat{\mu},\mu}^{-1}\psi_{x-\widehat{\mu}} - 2\psi_x \right). \tag{5}$$

Then $D_{\mu,U^g}^+(g\psi) = g D_{\mu,U}^+ \psi$ and $\Delta_{U^g}(g\psi) = g\Delta_U \psi$.

Proof. For the positive edge, substitute (2). The two gauge factors at $x + \hat{\mu}$ cancel, leaving g_x on the left. For the negative edge,

$$(U_{x-\widehat{\mu},\mu}^g)^{-1} g_{x-\widehat{\mu}} = g_x U_{x-\widehat{\mu},\mu}^{-1}.$$

Every term in either difference therefore has the common factor g_x. □

The lattice inner product is $\langle \phi, \psi \rangle_V = \sum_x \phi_x^* \psi_x$, conjugate-linear in the first entry. Periodic reindexing and unitarity also give the useful identity

$$-\langle \psi, \Delta_U \psi \rangle_V = \sum_{x,\mu} \left\| (D_{\mu,U}^+ \psi)_x \right\|^2. \tag{6}$$

Indeed, expanding one squared difference gives $h^{-2}(\|\psi_{x+\widehat{\mu}}\|^2 + \|\psi_x\|^2 - 2\,\mathrm{Re}(\psi_x^* U_{x,\mu}\psi_{x+\widehat{\mu}}))$. Summing and shifting the first norm term produces the left side. The forward and backward hopping blocks are adjoints, so Δ_U is self-adjoint and nonpositive. This finite identity uses the periodic boundary condition; an open domain needs its boundary terms or a specified boundary condition.

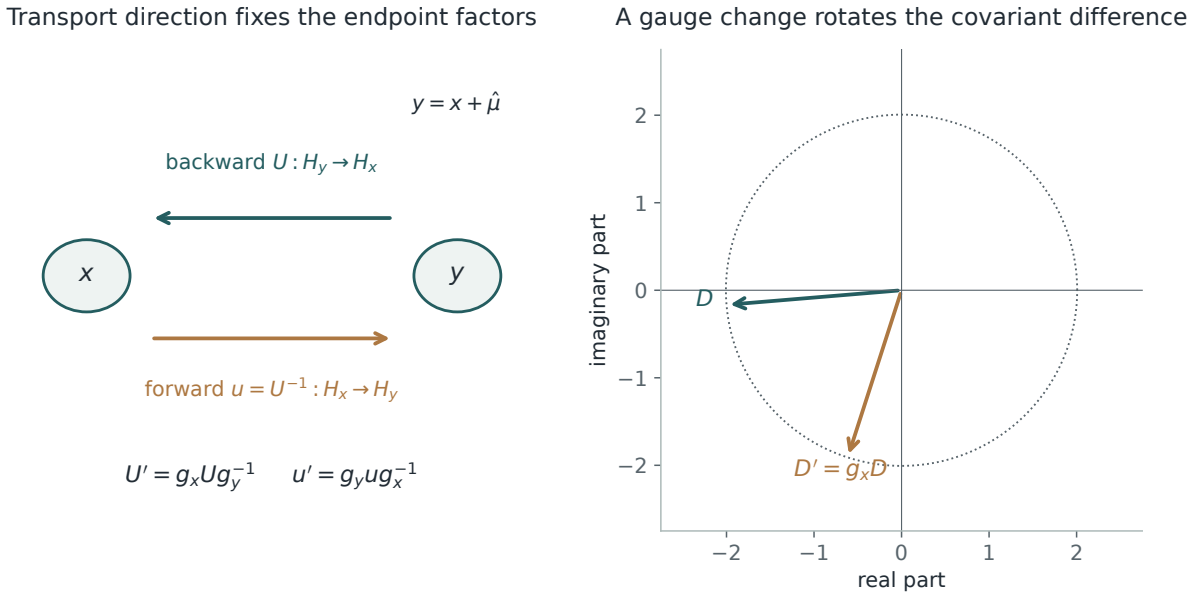


Figure 1. A backward link U and its forward inverse u compare the same two fibers. Gauge factors cancel at the input and remain at the destination. The scalar example illustrates the exact transformation law.

2 A linked walk and its actual charge flux

Covariance alone does not specify a current. To obtain one, fix the walk that moves the amplitudes. Let $N \geq 1$ be an integer, let $j \in \mathbb{Z}/N\mathbb{Z}$, let $\psi_j^n \in \mathbb{C}^2$, and let $U_j^n \in U(1)$ be backward spatial links. Let $T_j^n \in U(1)$ transport from (j, n) to $(j, n+1)$. At each step choose two on-site unitary coins $C_{1,j}^n, C_{2,j}^n \in U(2)$ and define

$$\eta_j = C_{1,j}^n \psi_j^n, \tag{7}$$

$$\xi_j^+ = (U_{j-1}^n)^{-1}\eta_{j-1}^+, \qquad \xi_j^- = U_j^n \eta_{j+1}^-, \tag{8}$$

$$\psi_j^{n+1} = T_j^n C_{2,j}^n \xi_j. \tag{9}$$

Here the $+$ channel moves toward increasing j, while the $-$ channel moves toward decreasing j. The legacy implementation takes the two coins to be equal half mass rotations. The argument needs only their unitarity and the fact that scalar phases commute with them.

For a time-dependent scalar gauge g_j^n, use

$$(U_j^n)^g = g_j^n U_j^n (g_{j+1}^n)^{-1}, \qquad (T_j^n)^g = g_j^{n+1} T_j^n (g_j^n)^{-1}. \tag{10}$$

Spatial transport first brings the gauge factor to site j on the old slice. Temporal transport then brings it to the new slice. Thus the step operator $\mathcal{W}_n$ satisfies

$$\mathcal{W}_n[U^g, T^g]G^n = G^{n+1}\mathcal{W}_n[U, T], \qquad (G^n\psi)_j = g_j^n\psi_j. \tag{11}$$

Every factor in the step is unitary on the finite state space: the two coins and temporal phases act sitewise, and the linked shift is a permutation of the separate channels followed by unit-modulus multiplication.

Proposition 2.1 (Exact one-step continuity)**.** *Set* $\rho_j^n = \left\|\psi_j^n\right\|^2$ *and*

$$J_{j+1/2}^n = |\eta_j^+|^2 - |\eta_{j+1}^-|^2. \tag{12}$$

The linked walk obeys

$$\rho_j^{n+1} - \rho_j^n = J_{j-1/2}^n - J_{j+1/2}^n. \tag{13}$$

Both the density and this current are gauge invariant. Total charge is conserved. For a consecutive periodic arc A *with specified left and right boundary edges,* $Q_A^{n+1} - Q_A^n = J_{\partial_- A}^n - J_{\partial_+ A}^n$.

Proof. The first coin preserves the local norm, so $\rho_j^n = |\eta_j^+|^2 + |\eta_j^-|^2$. The second coin and temporal phase preserve the norm after the shift, giving $\rho_j^{n+1} = |\eta_{j-1}^+|^2 + |\eta_{j+1}^-|^2$. Subtract these expressions and group the two terms associated with each boundary edge. This is (13). The transformed η_j is $g_j^n\eta_j$, so its channel norms are unchanged. Summing the balance over the full cycle or the chosen arc cancels its internal currents. □

This current measures charge transferred during one discrete step. Dividing it by a chosen time step would give a rate. It is not the same expression as the instantaneous Hamiltonian current introduced below. The two expressions belong to different evolution laws, even though both produce a local balance by cancellation.

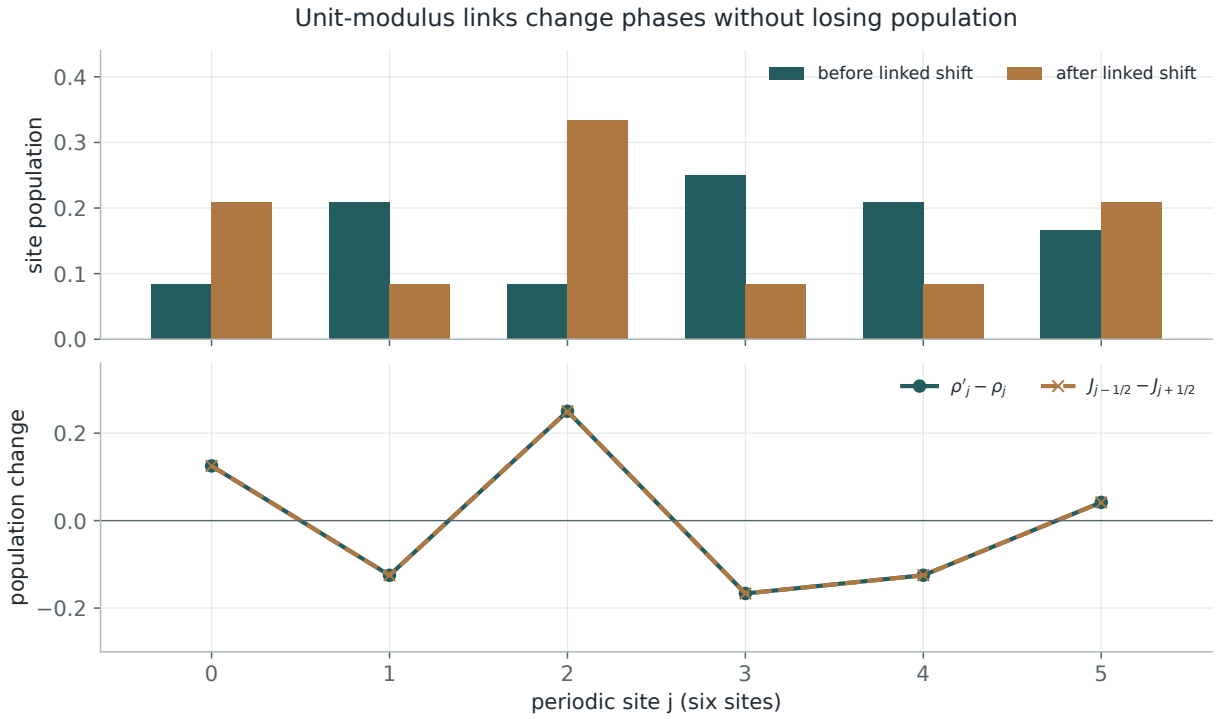


Figure 2. The explicit finite coined walk supplies a current before a continuum interpretation is attempted. The plotted computation checks one-step local balance and total population for its declared finite data.

3 Unit-circle Clifford channels in the formal transport layer

The later formal development makes the directional channels depend on the Clifford matrix of that direction. Let $A = A^*$ and $A^2 = I$ on a finite spinor space. Then

$$P_A^+ = \tfrac{1}{2}(I + A), \qquad P_A^- = \tfrac{1}{2}(I - A) \tag{14}$$

are self-adjoint complementary orthogonal projections. Idempotence and $P_A^+ P_A^- = 0$ follow by expanding the products and using $A^2 = I$; self-adjointness follows from $A = A^*$. Different spatial Clifford matrices generally have different spectral projections. A permanent assignment of the four spinor coordinates to two Boolean channels cannot replace this construction in all three directions.

For a bijection $\tau : V \to V$ and a forward scalar link $u : V \to U(1)$, define

$$(S_{A,\tau,u}\psi)(x) = u(\tau^{-1}x)P_A^+\psi(\tau^{-1}x) + u(x)^{-1}P_A^-\psi(\tau x). \tag{15}$$

This agrees with the orientation dictionary (3). Its inverse is

$$(S_{A,\tau,u}^{-1}\phi)(x) = u(x)^{-1}P_A^+\phi(\tau x) + u(\tau^{-1}x)P_A^-\phi(\tau^{-1}x). \tag{16}$$

Substitution proves both inverse identities: same-channel terms cancel the link with its inverse and mixed-channel terms vanish. Orthogonality, followed by reindexing x through τ in each channel, gives

$$\sum_x \langle S\psi(x), S\phi(x)\rangle = \sum_x \langle \psi(x), \phi(x)\rangle \,. \tag{17}$$

For $u^g(x) = g(\tau x)u(x)g(x)^{-1}$, direct endpoint cancellation gives $S_{A,\tau,u^g}G_g = G_g S_{A,\tau,u}$. Products of these directional steps and scalar-gauge-equivariant unitary mass steps retain covariance and inner-product preservation.

The genuine unit-circle hypothesis is essential. A nonzero complex scalar is invertible, but multiplication by 2 multiplies squared norm by 4. The phase must therefore lie in the unit circle, not merely in $\mathbb{C}^\times$. Norm preservation depends on that mathematical hypothesis. The formalization companion identifies its precise type and transport conventions.

The checked four-spinor construction in this layer uses scalar $U(1)$ phases. The $SU(2)$ matrices of the legacy non-Abelian chapter act on an internal two-dimensional fiber. To combine that color action with a Dirac spinor, one would specify a space such as $\mathbb{C}^4 \otimes \mathbb{C}^2$, with spin projectors on the first factor and gauge transport on the second. This guarantees their commutation. No assertion here identifies the two-dimensional color fiber with the four-dimensional Clifford spinor or declares the resulting non-Abelian Dirac extension checked.

4 Plaquettes, invariant observables, and curvature phases

On a two-dimensional coordinate plane, the backward links define the based plaquette

$$W_{x,\mu\nu} = U_{x,\mu}U_{x+\widehat{\mu},\nu}U^{-1}_{x+\widehat{\nu},\mu}U^{-1}_{x,\nu}. \tag{18}$$

Starting with a vector at x, the rightmost factor moves it to $x+\widehat{\nu}$; the next moves it to the far corner; the last two bring it back to x. The matrix product is therefore an endomorphism of the base fiber. Each intermediate gauge factor cancels in its written order, giving

$$W^g_{x,\mu\nu} = g_x W_{x,\mu\nu} g_x^{-1}. \tag{19}$$

For $U(1)$ the plaquette itself is invariant. For a noncommutative group it is its conjugacy class that is invariant; the matrix entries depend on the frame at x.

For $SU(2)$ links, the Wilson action of the source is the finite function

$$S_W[U] = \frac{1}{g^2_{\mathrm{YM}}} \sum_{x,\mu<\nu} \left(1 - \tfrac{1}{2}\operatorname{Re}\operatorname{Tr} W_{x,\mu\nu}\right), \qquad g_{\mathrm{YM}} \in \mathbb{R}\setminus\{0\}. \tag{20}$$

The subscript distinguishes the coupling from the site gauge g_x. Cyclicity of trace proves $S_W[U^g] = S_W[U]$. Since an $SU(2)$ matrix has eigenvalues $e^{i\theta}, e^{-i\theta}$, a plaquette term is $1-\cos\theta \geq 0$. In the Abelian case the analogous term is $1 - \operatorname{Re} W$. Specifying an invariant action is not yet a construction of its Hamiltonian phase space, Gauss constraint, or a continuum interacting field theory.

4.1 An Abelian surface product

For scalar links, multiply the oriented plaquettes on the six faces of a cube. Every edge occurs twice, with opposite orientation. Commutativity permits its factors to meet and cancel, so

$$\prod_{f\subset\partial c} W_f^{\mathrm{or}(f,c)} = 1. \tag{21}$$

The same cancellation makes the product of fine plaquettes on a rectangular patch equal its boundary holonomy. The statement is an exact group identity. If $W_f = e^{i\theta_f}$ with each θ_f chosen as a principal

argument, it implies that the oriented phase sum lies in $2\pi\mathbb{Z}$. It does not force that sum to be zero as a real number. An unwrapped small-phase regime or a specified branch is needed for that stronger conclusion.

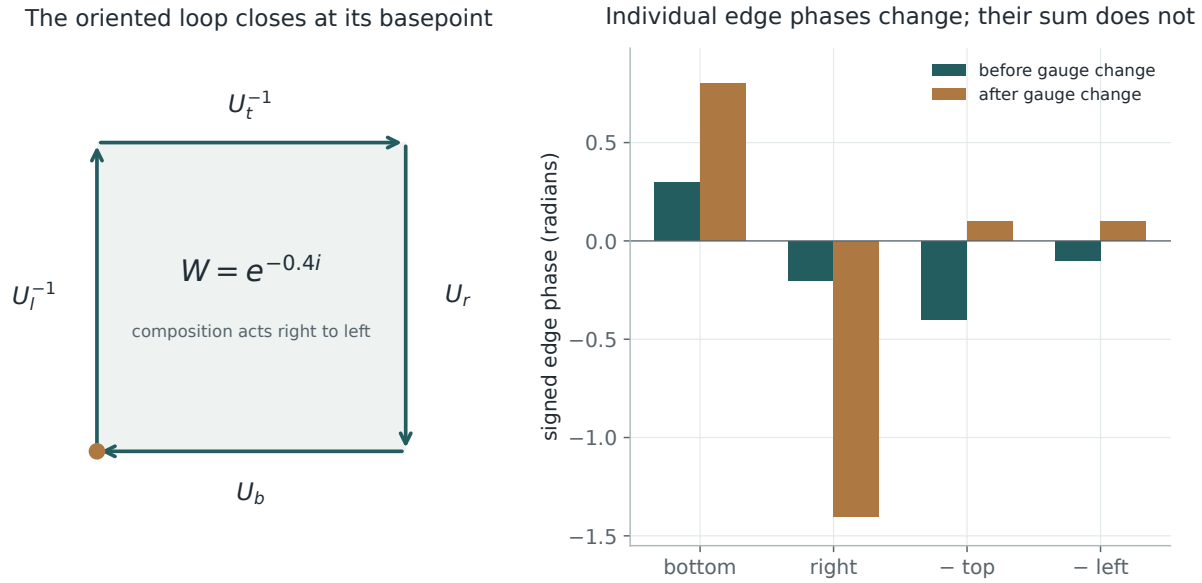


Figure 3. For Abelian links, the oriented edge phases cancel around adjacent faces. A holonomy is branch independent; a displayed real phase requires a branch choice.

4.2 Small smooth plaquettes

Suppose a smooth Abelian connection is sampled by exact line integrals, with a consistent loop orientation. Write its holonomy as $W = e^{i\theta}$, where the small real phase is $\theta = -q\int_{\square} F_{\mu\nu}\, dS$. Here q is the charge and $F_{\mu\nu}$ is the corresponding oriented curvature component. Take rectangles with side lengths $a_\mu, a_\nu = O(a)$ and uniformly controlled derivatives of F. If the rectangle is centered at x_c and F is C^2, symmetric integration removes its linear Taylor terms:

$$\theta = -q a_\mu a_\nu F_{\mu\nu}(x_c) + O(a^4). \tag{22}$$

Consequently,

$$1 - \operatorname{Re} W = \tfrac{1}{2} q^2 a_\mu^2 a_\nu^2 F_{\mu\nu}(x_c)^2 + O(a^6). \tag{23}$$

To see the remainder, use $1-\cos\theta = \theta^2/2 + O(\theta^4)$ and $\theta = O(a^2)$. The cross term from the $O(a^4)$ phase error is $O(a^6)$, while the fourth-power term is $O(a^8)$. Evaluating the curvature at a corner instead generally leaves an $O(a^3)$ phase error and hence an $O(a^5)$ action-density error. These orders concern the chosen pointwise comparison and exact-integral sampling. They do not by themselves justify convergence of an action sum or its solutions.

5 A finite Maxwell–Dirac model and its conserved constraint

The finite lattice can also carry dynamical link variables. The following is the particular semidiscrete model in the Abelian source, with time continuous and space a periodic cycle of $N \geq 3$ sites. It is useful because its matter current, electric work, and Gauss propagation can all be checked in the same equations.

Fix $h > 0$, $m \in \mathbb{R}$, and Hermitian spin matrices $\alpha = \alpha^*$ and $\beta = \beta^*$ on $\mathbb{C}^s$. Let $\psi_j \in \mathbb{C}^s$ and $A_j, E_j \in \mathbb{R}$, and set $U_j = e^{-ihA_j}$. The source absorbs the charge into A. Its real Hamiltonian is

$$H_h = H_{\mathrm{m}} + H_E, \tag{24}$$

$$H_{\mathrm{m}} = hm \sum_j \psi_j^* \beta \psi_j + \sum_j \mathrm{Im}(\psi_j^* \alpha U_j \psi_{j+1}), \qquad H_E = \frac{h}{2} \sum_j E_j^2. \tag{25}$$

Hermiticity is enough for the identities below. Additional Clifford relations specify a Dirac representation but are not needed to prove charge or energy conservation in this finite model. The action fixes all normalization and sign conventions:

$$\mathcal{S}_h = \int \left[\frac{ih}{2} \sum_j (\psi_j^* \dot{\psi}_j - \dot{\psi}_j^* \psi_j) + h \sum_j E_j \dot{A}_j - H_h \right] dt. \tag{26}$$

Variations with fixed endpoint data yield $ih\dot{\psi}_j = \partial H_h / \partial \overline{\psi}_j$, $h\dot{A}_j = \partial H_h / \partial E_j$, and $h\dot{E}_j = -\partial H_h / \partial A_j$. Written out, these are

$$\dot{\psi}_j = -im\beta\psi_j - \frac{1}{2h} \alpha U_j \psi_{j+1} + \frac{1}{2h} U_{j-1}^* \alpha \psi_{j-1}, \tag{27}$$

$$\dot{A}_j = E_j, \qquad \dot{E}_j = J_j, \qquad J_j = \mathrm{Re}(\psi_j^* \alpha U_j \psi_{j+1}). \tag{28}$$

For example, if $z_j = \psi_j^* \alpha U_j \psi_{j+1}$, then $\partial_{A_j} \operatorname{Im} z_j = \mathrm{Im}(-ihz_j) = -h \operatorname{Re} z_j$. This fixes the positive sign of $\dot{E}_j = J_j$ for the conventions in the action.

Theorem 5.1 (Charge balance and Gauss propagation)**.** *For every differentiable solution of* (27)–(28)*, define*

$$\rho_j = \psi_j^* \psi_j, \qquad \bar{\rho} = \frac{1}{N} \sum_j \rho_j, \qquad \mathcal{G}_j = \frac{E_j - E_{j-1}}{h} + \rho_j - \bar{\rho}. \tag{29}$$

Then

$$\dot{\rho}_j + \frac{J_j - J_{j-1}}{h} = 0, \qquad \dot{\bar{\rho}} = 0, \qquad \dot{\mathcal{G}}_j = 0. \tag{30}$$

In particular, initial data satisfying $\mathcal{G}_j = 0$ retain that constraint on their interval of existence.

Proof. Differentiate ρ_j and use $\dot{\rho}_j = 2\,\mathrm{Re}(\psi_j^* \dot{\psi}_j)$. The mass term contributes zero, since $\psi_j^* \beta \psi_j$ is real. The forward hopping term contributes $-J_j/h$. The backward scalar is the conjugate of $\psi_{j-1}^* \alpha U_{j-1} \psi_j$, using $\alpha = \alpha^*$ and the scalar nature of U. Its contribution is J_{j-1}/h. This proves continuity. Periodic summation gives $\dot{\bar{\rho}} = 0$. Finally,

$$\dot{\mathcal{G}}_j = \frac{\dot{E}_j - \dot{E}_{j-1}}{h} + \dot{\rho}_j - \dot{\bar{\rho}} = \frac{J_j - J_{j-1}}{h} - \frac{J_j - J_{j-1}}{h} = 0.$$

□

The subtraction of $\bar{\rho}$ is structural on a periodic cycle: the discrete divergence sums to zero, so a constraint containing a nonzero total charge without compensating background would be inconsistent. Equation (29) supplies that neutralizing mean. Conservation of $\mathcal{G}_j$ is a propagation theorem. It preserves an initially nonzero constraint violation too; it does not project arbitrary initial data onto the constraint surface. Summing continuity over $a \le j \le b$ gives the physical-length-weighted charge balance

$$\frac{d}{dt}\left(h\sum_{j=a}^{b}\rho_j\right) = J_{a-1} - J_b.$$

Proposition 5.2 (Electric work and total energy)**.** *Along the same solutions,*

$$\dot{H}_E = h\sum_j E_j J_j, \qquad \dot{H}_{\mathrm{m}} = -h\sum_j E_j J_j, \qquad \dot{H}_h = 0. \tag{31}$$

Proof. The first formula follows from $\dot{E}_j = J_j$. For the matter contribution, let $q_j = \partial H_{\mathrm{m}}/\partial\overline{\psi}_j = ih\dot{\psi}_j$. The state-dependent part of the real derivative is $2\,\mathrm{Re}\sum_j q_j^*\dot{\psi}_j = 0$, because $q_j^*\dot{\psi}_j = -ih\left\|\dot{\psi}_j\right\|^2$ is purely imaginary. The remaining explicit link derivative is

$$\sum_j \frac{\partial H_{\mathrm{m}}}{\partial A_j}\dot{A}_j = -h\sum_j J_j E_j.$$

Adding the two sectors proves total conservation. □

5.1 Which gauge transformations remain in this action?

For time-independent real χ_j, the transformation

$$\psi_j \mapsto e^{i\chi_j}\psi_j, \qquad A_j \mapsto A_j + \frac{\chi_{j+1}-\chi_j}{h}, \qquad E_j \mapsto E_j \tag{32}$$

leaves the action and equations invariant. In fact U_j then has precisely the backward-link law (2), every hopping scalar z_j is unchanged, and the time derivatives of the gauge parameters vanish.

This is the residual symmetry of the stated temporal-gauge model. If χ_j depends on time, the matter time derivative acquires $i\dot{\chi}_j\psi_j$, while $\dot{A}_j$ acquires $(\dot{\chi}_{j+1}-\dot{\chi}_j)/h$. With E_j fixed and no scalar potential in (26), these changes do not preserve the displayed equations. The fully time-dependent covariance of the discrete linked walk uses its temporal link T; it must not be transferred to this action after that degree of freedom has been removed.

The statements in this section are ordinary proofs for a declared finite-dimensional classical field model. They establish genuine interaction between its matter and electric variables. They do not constitute a Lean construction of a coupled gauge Hamiltonian, a quantized gauge field, or a global continuum Maxwell–Dirac solution.

6 What a second-order covariant stencil requires

The exact finite identities did not assume a smooth field in the background. A continuum comparison does. Let A be a real C^2 function near x, let ψ be C^3 there, and set

$$U_h(x) = \exp\left(-i\int_x^{x+h} A(s)\,ds\right), \qquad (D_h^A\psi)(x) = \frac{U_h(x)\psi(x+h) - U_h(x-h)^{-1}\psi(x-h)}{2h}. \tag{33}$$

The smooth A in this paragraph is a continuum coefficient. In the finite Hamiltonian one obtains the same link by taking A_j to be its edge average $h^{-1}\int_{x_j}^{x_j+h} A(s)\,ds$. This is an additional sampling prescription.

Proposition 6.1 (Local second-order consistency)**.** *With uniform bounds on A and its derivatives through order two and on ψ and its derivatives through order three on a fixed neighborhood,*

$$(D_h^A\psi)(x) = (\partial_x - iA(x))\psi(x) + O(h^2) \qquad (h\to 0). \tag{34}$$

The remainder is uniform over points whose neighborhoods have the same derivative bounds.

Proof. For fixed x, put $f_x(s) = \exp(-i\int_x^{x+s} A(r)\,dr)\psi(x+s)$. Then the two numerator terms in (33) are $f_x(h)$ and $f_x(-h)$. The derivative at zero is $f_x'(0) = \psi'(x) - iA(x)\psi(x)$. Taylor expansion of f_x at zero to second order with bounded third derivative gives $f_x(h) - f_x(-h) = 2hf_x'(0) + O(h^3)$. Division by $2h$ proves the assertion. The derivative bounds for A and ψ give the required uniform bound on f_x'''. □

A midpoint link $\widetilde{U}_h(x) = e^{-ihA(x+h/2)}$ has the same conclusion under these bounds. The midpoint quadrature error in the integral is $O(h^3)$, and $|e^{ia} - e^{ib}| \le |a-b|$ for real a, b transfers that error to the link. The two numerator perturbations are each $O(h^3)$ for bounded ψ, hence the stencil perturbation is $O(h^2)$.

Example 6.2 (Left-endpoint sampling can lose an order)**.** Take $A(x) = x$ near the origin and $\psi(x) = 1$. With the left-endpoint prescription $\widehat{U}_h(x) = e^{-ihA(x)}$, the stencil at zero is

$$\frac{\widehat{U}_h(0) - \widehat{U}_h(-h)^{-1}}{2h} = \frac{1 - e^{-ih^2}}{2h} = \frac{i}{2}h + O(h^3).$$

The continuum value $(\partial_x - iA(0))\psi(0)$ is zero. Thus this perfectly smooth input has a first-order error. The formula $U_j = e^{-ihA_j}$ by itself does not identify a second-order sampling scheme.

This qualification does not alter the finite Hamiltonian or its conservation laws. It isolates what is needed to interpret its link values as samples of a continuum connection. Nor does consistency of one stencil prove convergence of the coupled time evolution: stability, compatible initial data, regularity over the comparison interval, and a choice of reconstruction still have to enter that argument.

7 Non-Abelian currents and their frame dependence

Now take $\psi_x \in \mathbb{C}^2$, $U_{x,\mu} \in SU(2)$, and local gauges $g_x \in SU(2)$, as in the legacy non-Abelian chapter. The linked scalar

$$z_{x,\mu} = \psi_x^* U_{x,\mu}\psi_{x+\widehat{\mu}} \tag{35}$$

is gauge invariant: the adjoint factor g_x^* cancels g_x, and the factors at $x + \hat{\mu}$ also cancel. It follows that the source's scalar link observable

$$j_{x,\mu} = 2\,\mathrm{Im}\, z_{x,\mu} \tag{36}$$

is invariant. To obtain a continuity equation with its particular sign and normalization one must specify the hopping generator. Invariance of the observable alone is not an equation of motion for color charge.

Let $\sigma^1, \sigma^2, \sigma^3$ be the Pauli matrices and define the real components $q_x^a = \psi_x^* \sigma^a \psi_x$. Expand the conjugated Pauli basis as

$$g_x^* \sigma^a g_x = \sum_b R_x{}^a{}_b \sigma^b, \qquad q_x^{g,a} = \sum_b R_x{}^a{}_b q_x^b. \tag{37}$$

The coefficients are real because conjugation preserves traceless Hermitian matrices. The trace inner product $\frac{1}{2}\,\mathrm{Tr}(AB)$ is preserved by conjugation, and the Pauli basis is orthonormal for it; hence $R_x R_x^{\mathsf{T}} = I$. Thus the color vector rotates while its Euclidean squared length stays fixed. The scalar charge $\psi_x^* \psi_x$ is separately invariant. A color component relative to a selected Pauli axis is not generally invariant under all local frames.

The same distinction applies to a plaquette: W changes by conjugation, while its trace does not. Keeping the full matrix retains transport information that a single trace can lose. Even for $SU(2)$, identical traces at separate plaquettes do not determine the relative frames needed to combine them into a larger loop.

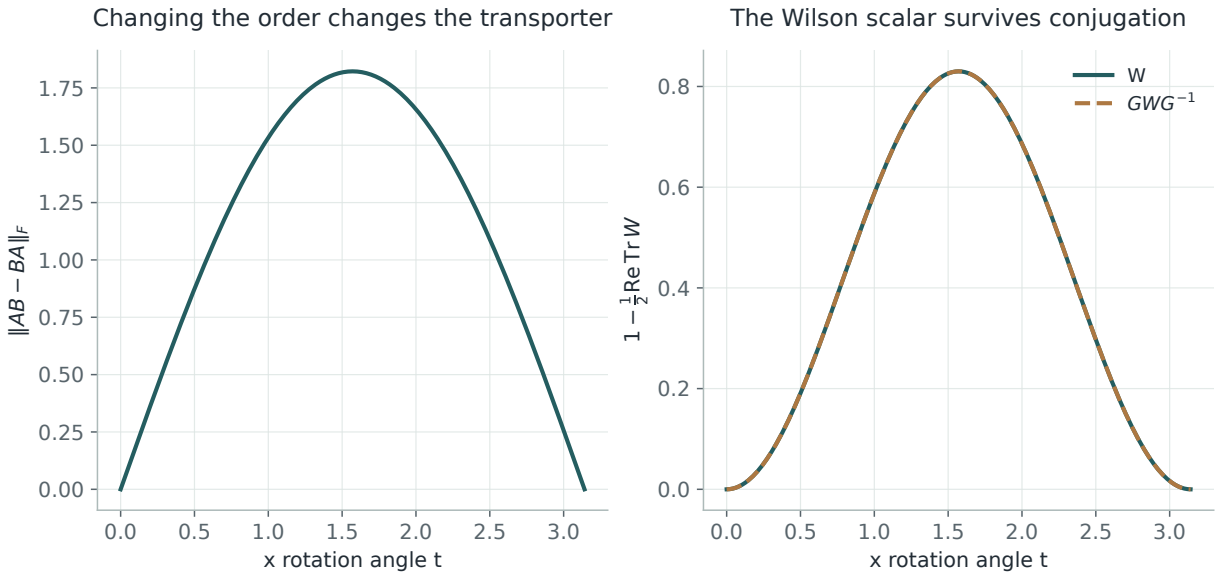


Figure 4. Noncommuting $SU(2)$ factors retain their order. A based holonomy changes by conjugation, while the normalized Wilson trace is unchanged. The color fiber here is separate from the four-component Clifford spinor.

8 Exact ordered refinement of a two-by-two patch

The Abelian rule "multiply the fine plaquettes" hides two choices in a noncommutative group: where each plaquette is based, and in which order the transported matrices are multiplied. A 2×2 patch makes both choices visible without an asymptotic expansion.

Write X_{ij} for the backward horizontal link from $(i+1, j)$ to (i, j), where $i = 0, 1$ and $j = 0, 1, 2$. Write Y_{ij} for the backward vertical link from $(i, j+1)$ to (i, j), where $i = 0, 1, 2$ and $j = 0, 1$. Define the four fine plaquettes by

$$W_{ij} = X_{ij} Y_{i+1,j} X_{i,j+1}^{-1} Y_{ij}^{-1}, \qquad i, j \in \{0, 1\}. \tag{38}$$

Each W_{ij} acts at its own lower-left basepoint (i, j). The two row holonomies, based at $(0, j)$, are

$$R_j = (X_{0j} W_{1j} X_{0j}^{-1}) W_{0j}, \qquad j = 0, 1. \tag{39}$$

The conjugation transports the right plaquette into the left base fiber before multiplying it with the left plaquette.

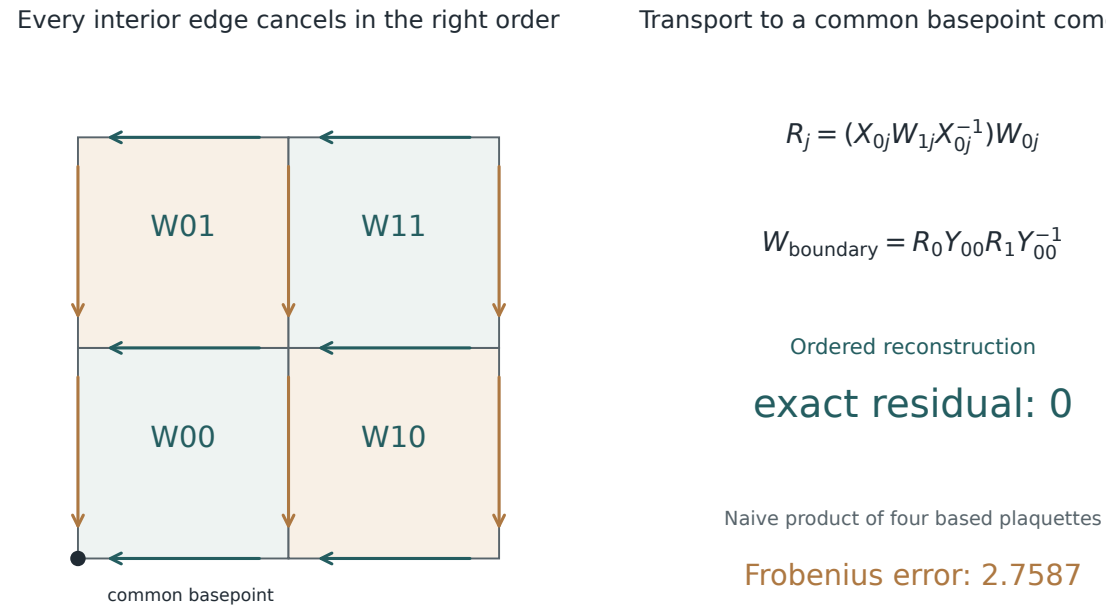


Figure 5. The four fine plaquettes are brought to common basepoints in a specified order. The exact cancellation leaves the eight outer links; a naive product omits the transport information.

Theorem 8.1 (Ordered patch identity). *For invertible links in any matrix group, the coarse boundary holonomy satisfies*

$$W^{(2)} := X_{00} X_{10} Y_{20} Y_{21} X_{12}^{-1} X_{02}^{-1} Y_{01}^{-1} Y_{00}^{-1} \tag{40}$$

$$= R_0 Y_{00} R_1 Y_{00}^{-1}. \tag{41}$$

This identity requires neither small links nor commuting links.

Proof. First expand one row while retaining the product order:

$$\begin{aligned} R_j &= X_{0j}X_{1j}Y_{2j}X_{1,j+1}^{-1}Y_{1j}^{-1}X_{0j}^{-1}X_{0j}Y_{1j}X_{0,j+1}^{-1}Y_{0j}^{-1} \\ &= X_{0j}X_{1j}Y_{2j}X_{1,j+1}^{-1}X_{0,j+1}^{-1}Y_{0j}^{-1}. \end{aligned}$$

The adjacent pairs $X_{0j}^{-1}X_{0j}$ and $Y_{1j}^{-1}Y_{1j}$ cancel. Substituting these row expressions gives

$$\begin{aligned} R_0Y_{00}R_1Y_{00}^{-1} = X_{00}X_{10}Y_{20}X_{11}^{-1}X_{01}^{-1}\underbrace{Y_{00}^{-1}Y_{00}}_{I} \\ \cdot X_{01}X_{11}Y_{21}X_{12}^{-1}X_{02}^{-1}Y_{01}^{-1}Y_{00}^{-1}. \end{aligned}$$

Now $X_{01}^{-1}X_{01}$ cancels, followed by $X_{11}^{-1}X_{11}$. The remaining factors are exactly (40). □

All terms in (41) transform by conjugation at $(0,0)$ once their displayed transports are included. This checks the basepoint typing as well as the algebra. In the Abelian case conjugations disappear and the fine plaquettes can be reordered, recovering the familiar surface product. In the non-Abelian case a naive product of the four untransported matrices can be formed after identifying their coordinate spaces, but it has no reason to obey the required local gauge transformation or to equal the boundary holonomy.

This refinement is exact because it refers to one fixed set of fine links and its outer boundary. It does not assert that independently assigned coarse links agree with products of fine links. That separate compatibility question is the starting point of the geometric refinement estimates in the next part of the monograph.

9 Extracting a local curvature and cancelling its first error

Exact group multiplication and approximate Lie-algebra curvature are different operations. For fixed $A, B \in \mathfrak{su}(2)$ define the commutator loop

$$W_h = e^{hA}e^{hB}e^{-hA}e^{-hB}. \tag{42}$$

Here A and B are constant skew-Hermitian traceless matrices, and h is a real scalar near zero. Because $W_0 = I$, a single analytic matrix logarithm near the identity can be used for all sufficiently small $|h|$. Define

$$F_h = h^{-2}\log W_h, \qquad h \neq 0. \tag{43}$$

This construction describes a local constant-matrix model of noncommuting transport. A varying connection has additional derivative terms and must be analyzed with those terms present.

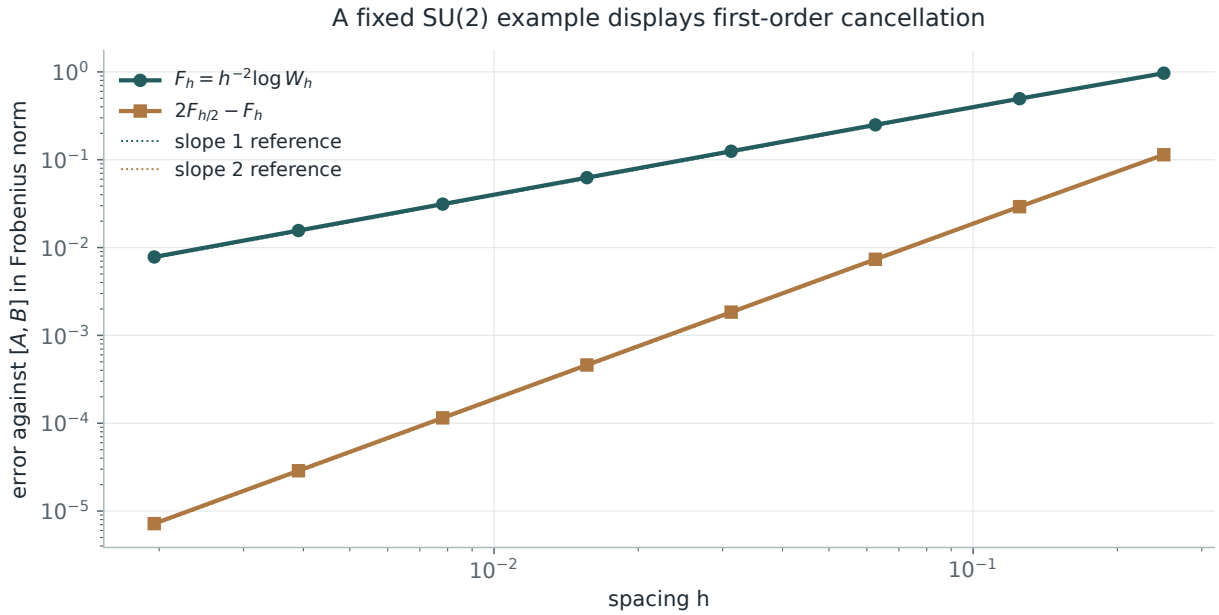


Figure 6. For the specified constant generators, the finite calculation illustrates first-order raw curvature error and cancellation of its leading term. The analytic statement is Proposition 9.1; a fitted slope over finitely many scales is a diagnostic of this example.

Proposition 9.1 (Local logarithmic refinement)**.** *On the common analytic logarithm branch, in any fixed matrix norm,*

$$F_h = [A, B] + hC_3(A, B) + O(h^2), \tag{44}$$

$$C_3(A, B) = \tfrac{1}{2}[A + B, [A, B]], \tag{45}$$

$$F_h^{[1]} := 2F_{h/2} - F_h = [A, B] + O(h^2). \tag{46}$$

For constant $G \in SU(2)$, simultaneous conjugation of A, B conjugates both F_h and $F_h^{[1]}$ by G.

Proof. Multiplying the four exponential power series through degree three gives

$$W_h = I + h^2[A, B] + \frac{h^3}{2}[A + B, [A, B]] + O(h^4).$$

The analytic logarithm has $\log(I + Z) = Z + O(\|Z\|^2)$ near zero. Since $Z = W_h - I = O(h^2)$, its correction begins at order h^4. Dividing by h^2 proves (44). At half spacing the linear term is $hC_3/2$, so $2F_{h/2} - F_h$ removes it exactly and leaves an $O(h^2)$ remainder. Finally, the power series for both the exponential and the local logarithm commute with conjugation. Applying them term by term proves the covariance statements. □

The common logarithm branch is essential: unrelated choices can introduce terms of size h^{-2} after rescaling and invalidate the cancellation. The Richardson value is a linear combination of Lie-algebra estimates; it is not asserted to be the logarithm obtained by the exact patch rule.

Geometry, defects, and reconstruction

10 Geometric half-edges and exact directional refinement

An exact change of resolution needs both a geometry and a comparison map. Replicating each coarse state into two spectator copies gives a useful algebraic test, but it does not make the translation visit a new point between neighboring sites. The geometric construction instead inserts an actual midpoint and transports the state to it.

We use the forward convention of the formal refinement modules. A link $u(x)$ transports from x to τx, so

$$u^g(x) = g(\tau x)u(x)g(x)^{-1}.$$

Thus $u(x) = U(x)^{-1}$ when $U(x)$ denotes the backward link used in the covariant difference stencil. Products act from right to left. All links in this part belong to $U(1)$; the unit modulus assumption is used in every Hilbert norm statement. Let $A = A^*$, $A^2 = 1$, act on $\mathbb{C}^4$, and put $P_\pm = (1 \pm A)/2$. The directional transport is

$$(S_{A,\tau,u}\psi)(x) = u(\tau^{-1}x)P_+\psi(\tau^{-1}x) + u(x)^{-1}P_-\psi(\tau x). \tag{47}$$

The channels are orthogonal. A bijective site translation and unit links therefore make this an isometry for the finite counting-space inner product established earlier.

10.1 Midpoint geometry and independent gauges

Write e_x for the retained endpoint and m_x for the midpoint of the edge $x \to \tau x$. The fine translation is the bijection

$$\sigma e_x = m_x, \qquad \sigma m_x = e_{\tau x}.$$

Consequently $\sigma^2 e_x = e_{\tau x}$ and $\sigma^2 m_x = m_{\tau x}$. This is the precise sense in which two half translations produce one coarse translation. Assigning half the edge spacing to σ is compatible with this combinatorics; it is not by itself a chart or continuum limit theorem.

Choose unit links $a(x)$ and $b(x)$ on the first and second half-edge with

$$b(x)a(x) = u(x). \tag{48}$$

The associated fine link is $v(e_x) = a(x)$, $v(m_x) = b(x)$. The interpolation and reconstruction maps are

$$(I_a\psi)(e_x) = \psi(x), \qquad (I_a\psi)(m_x) = a(x)\psi(x), \tag{49}$$

$$(R_a\Psi)(x) = \tfrac{1}{2}(\Psi(e_x) + a(x)^{-1}\Psi(m_x)). \tag{50}$$

The midpoint is a transported value, not an independently sampled field value. This identifies exactly the subspace on which the refinement theorem acts.

Coarse endpoint gauges $g(x)$ and midpoint gauges $q(x)$ may be chosen independently. The two sublinks transform by

$$a^g(x) = q(x)a(x)g(x)^{-1}, \qquad b^g(x) = g(\tau x)b(x)q(x)^{-1}.$$

The midpoint factors cancel in $b^g a^g$, leaving the required coarse transformation. At the midpoint,

$$a^g(x)g(x)\psi(x) = q(x)a(x)\psi(x),$$

while endpoint covariance is immediate. Hence interpolation commutes exactly with the corresponding gauge actions. This argument allows independent new gauge data at every inserted midpoint.

Theorem 10.1 (Exact two-half-step refinement)**.** *For every coarse state ψ and subdivision satisfying* (48)*,*

$$S^2_{A,\sigma,v} I_a \psi = I_a S_{A,\tau,u} \psi. \tag{51}$$

No flatness hypothesis is needed.

Proof. When the fine shift is applied twice, the mixed projector terms vanish because $P_+P_- = P_-P_+ = 0$. The remaining positive channel travels two half-edges forward, and the negative channel travels the same pair backward. At e_x, the positive coefficient is $b(\tau^{-1}x)a(\tau^{-1}x) = u(\tau^{-1}x)$; the negative coefficient is $a(x)^{-1}b(x)^{-1} = u(x)^{-1}$. The result is exactly (47). At m_x, the two surviving terms are

$$a(x)u(\tau^{-1}x)P_+\psi(\tau^{-1}x) + b(x)^{-1}P_-\psi(\tau x).$$

Since $b(x)^{-1} = a(x)u(x)^{-1}$, this is $a(x)$ times the coarse shifted state, as required by (49). □

10.2 The measure is part of the comparison

For finite sites, give each fine point half the coarse weight:

$$\langle \Psi, \Phi \rangle_{1/2} = \tfrac{1}{2} \sum_x \big(\langle \Psi(e_x), \Phi(e_x) \rangle + \langle \Psi(m_x), \Phi(m_x) \rangle \big).$$

Unitarity of $a(x)$ then gives

$$\langle I_a\psi, I_a\phi \rangle_{1/2} = \langle \psi, \phi \rangle , \qquad \langle I_a\psi, \Psi \rangle_{1/2} = \langle \psi, R_a\Psi \rangle . \tag{52}$$

The second identity follows by moving $a(x)$ from the first argument of the inner product to the second as $a(x)^{-1}$. It explains the averaging formula, including its factor $1/2$. In particular $R_aI_a = 1$ and I_aR_a is the orthogonal projection onto the transported midpoint subspace. Endpoint evaluation also gives a left inverse, but does not have this adjoint interpretation.

A local mass step can be refined factor by factor. If a local linear map N satisfies $N^2 = M$, then N commutes with the scalar interpolation. Consequently

$$N^2 S^2_{A,\sigma,v} N^2 I_a = I_a M S_{A,\tau,u} M.$$

Unitary spectral phases provide explicit choices of N. The refined expression replaces each coarse factor by its two half factors. It does not assert that two repetitions of a complete fine $NS_{A,\sigma,v}N$ walk equal the coarse walk: the intervening order would be different.

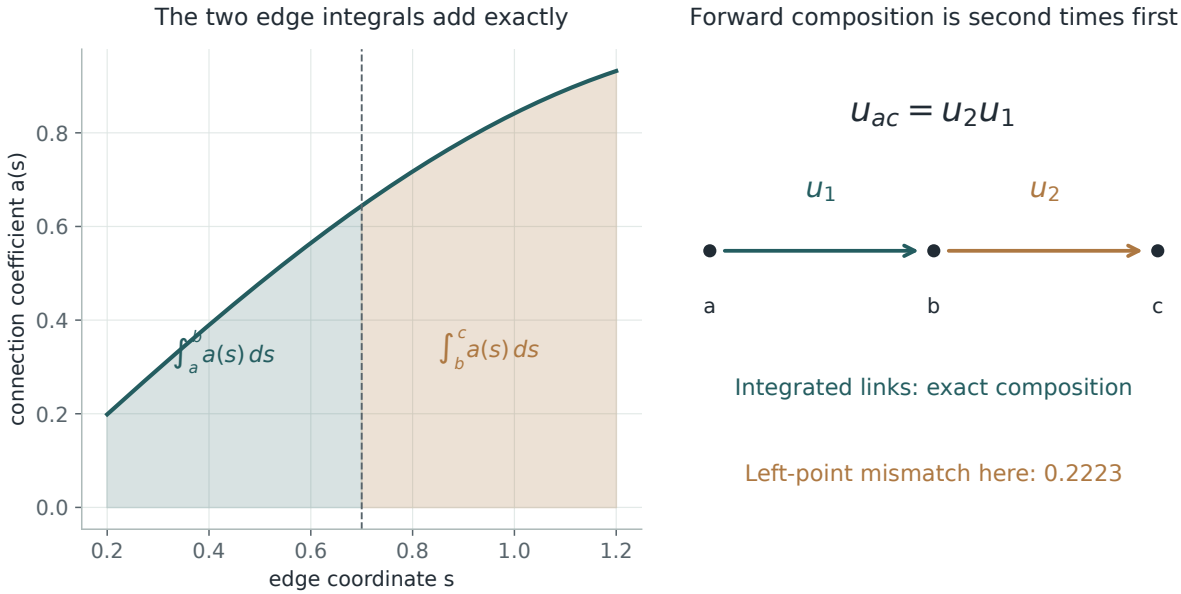


Figure 7. Integrated abelian phases compose exactly under edge subdivision. For the displayed coefficient $a(s) = \sin s$, the two subinterval integrals sum to the whole-edge integral. Independently sampled coarse and fine phases can disagree; that approximation error is separate from exact transport composition.

11 Cubical transport and the curvature obstruction

Three one-axis subdivisions do not automatically define a path-independent value at a shared corner. The issue already appears in one $2\times2\times2$ vertex cube. Let the vertices be $(x, y, z) \in \{0, 1\}^3$, with positive links $\ell_x(y, z)$, $\ell_y(x, z)$ and $\ell_z(x, y)$. There are twelve links and six faces. Under a vertex gauge, every link acquires its target gauge on the left and its source inverse on the right.

For example, the two positive paths across an xy face at height z are

$$p_{xy}(z) = \ell_y(1, z)\ell_x(0, z), \qquad p_{yx}(z) = \ell_x(1, z)\ell_y(0, z).$$

Define $\Omega_{xy}(z) = p_{xy}(z)p_{yx}(z)^{-1}$, and define the xz and yz face phases analogously. Both paths have the same endpoints, so their gauge factors cancel in this abelian ratio. Thus each Ω is gauge invariant, and $\Omega = 1$ precisely when the two path phases agree.

11.1 Six paths and one Bianchi relation

The three coordinate directions can be traversed in six orders, denoted $p_{XYZ}, p_{XZY}, p_{YXZ}, p_{YZX}, p_{ZXY}$ and p_{ZYX}. These names record execution order, while link multiplication still acts from right to left. Adjacent transpositions give

$$\begin{aligned} p_{XYZ} &= \Omega_{xy}(0)p_{YXZ}, & p_{XYZ} &= \Omega_{yz}(1)p_{XZY}, \\ p_{YXZ} &= \Omega_{xz}(1)p_{YZX}, & p_{XZY} &= \Omega_{xz}(0)p_{ZXY}, \\ p_{YZX} &= \Omega_{yz}(0)p_{ZYX}, & p_{ZXY} &= \Omega_{xy}(1)p_{ZYX}. \end{aligned}$$

Each identity follows by replacing the two-edge portion being swapped by its alternative path and its face ratio. The scalar nature of the links permits the common third edge to commute past the ratio.

Proposition 11.1 (Abelian cubical Bianchi identity)**.** *The six face phases satisfy*

$$\Omega_{yz}(1)\Omega_{xz}(0)\Omega_{xy}(1) = \Omega_{xy}(0)\Omega_{xz}(1)\Omega_{yz}(0). \tag{53}$$

If all six faces are flat, the six corner path phases agree. Flatness of the five faces other than $xy(1)$ already forces the sixth to be flat.

Proof. Use the three identities down the left route from XYZ through YXZ, YZX to ZYX, and the three down the right route through XZY, ZXY. They express p_{XYZ} as the respective products in (53) times p_{ZYX}. Cancel the unit terminal phase. If the five displayed independent face phases equal one, the identity gives $\Omega_{xy}(1) = 1$. When all faces are flat, every adjacent path swap leaves the phase unchanged, and the displayed relations connect all six paths. □

This is a local abelian identity. A matrix-valued version must account for the basepoint and transport of each face holonomy; the preceding formula cannot simply be copied with noncommuting matrices. Local flatness also says nothing by itself about holonomy around noncontractible loops in a larger lattice.

11.2 Interpolation is possible even when it depends on a path

Choose one reference path from the origin to each vertex and let $r(v)$ be its transport. Then

$$(I_\ell\psi)(v) = r(v)\psi, \qquad \langle\Psi,\Phi\rangle_{\text{cube}} = \tfrac{1}{8}\sum_{v\in\{0,1\}^3}\langle\Psi(v),\Phi(v)\rangle\,.$$

Every $r(v)$ is unitary, so this map is a Hilbert isometry from one spinor to the eight-child cell. Its gauge covariance follows from telescoping along the selected paths. Neither property requires flatness. What curvature obstructs is independence from the reference-path choice.

Suppose two such paths have phases p and q with $p = \Omega q$. Their discrepancy has the exact form

$$D(p,q;\psi) = p\psi - q\psi = q(\Omega-1)\psi, \qquad \|D(p,q;\psi)\|^2 = |\Omega-1|^2\,\|\psi\|^2\,. \tag{54}$$

This is an equality, not a small-curvature approximation. It follows by factoring q and using $|q| = 1$. For a nonzero test spinor, the defect vanishes exactly when $\Omega = 1$. The formal nonflat cube has $p_{XYZ} = -1$ and $p_{YXZ} = 1$, producing a nonzero defect on the first basis spinor. Thus the distinction between covariant interpolation and path-independent interpolation is witnessed by explicit data.

For later estimates, retain the five-face quantity

$$\begin{aligned}\kappa(\ell) = {}& |\Omega_{xy}(0)-1| + |\Omega_{xz}(1)-1| + |\Omega_{yz}(0)-1| \\ & + |\Omega_{yz}(1)-1| + |\Omega_{xz}(0)-1|.\end{aligned} \tag{55}$$

The omitted face is $xy(1)$. This selection is the one used by the formal five-swap composition; the multiplicities in later constants come from this choice.

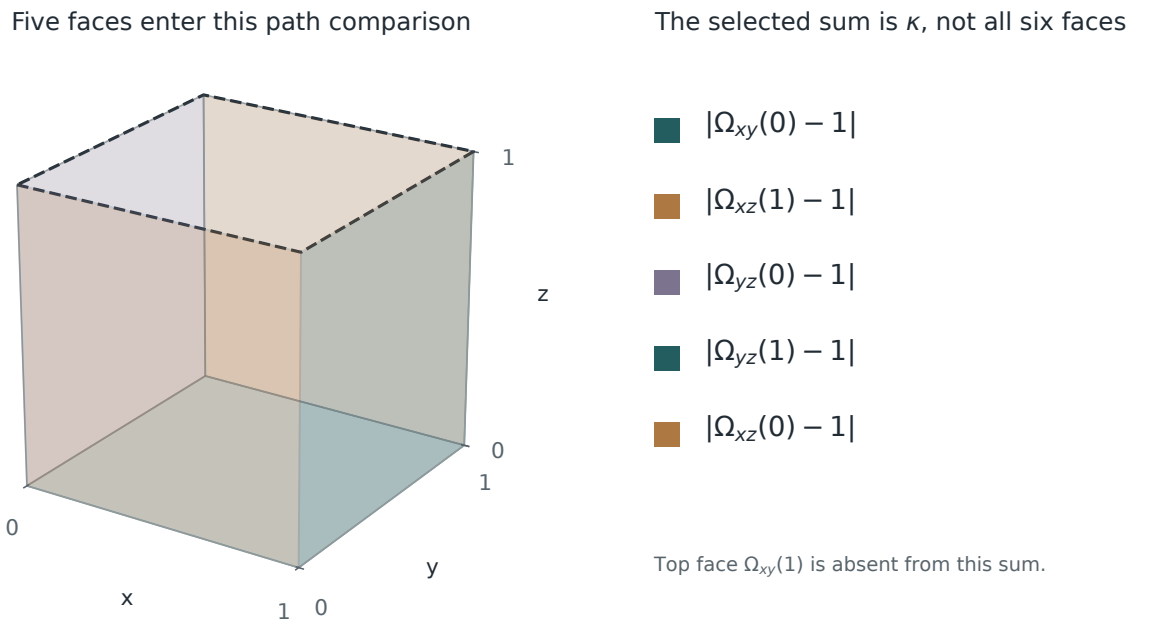


Figure 8. The selected five faces correspond to five adjacent swaps among the six corner paths. The sixth face is linked to them by the abelian cube Bianchi identity; it is not an independently omitted physical boundary condition.

12 Exact straight refinement and curvature are independent

The one-edge identity can be assembled into a three-axis refinement on the fine site set Site $\times \{0,1\}^3$. The half translation σ_i toggles the corresponding child bit; when that bit carries from one to zero, it applies the coarse translation τ_i to the site. Thus σ_i^2 is the lift of τ_i. A phase frame $r(s,c) \in U(1)$ defines

$$(I_r\psi)(s,c) = r(s,c)\psi(s).$$

With one-eighth fine weights this is again an isometry.

12.1 The straight compatibility contract

For each direction, canonical fine links can be defined by

$$v_i(p) = r(\sigma_i p)\,\gamma_i(p)\,r(p)^{-1},$$

where $\gamma_i(p) = 1$ on a noncarrying half step and $\gamma_i(p) = u_i(s)$ on a carrying half step. Cancellation of the intermediate frame gives

$$v_i(\sigma_i p)v_i(p) = r(\sigma_i^2 p)u_i(s)r(p)^{-1}. \tag{56}$$

For arbitrary fine links, equation (56) is the explicit straight compatibility hypothesis. It is gauge covariant, and it is exactly the hypothesis needed to repeat the projector argument of (51):

$$S_{i,\mathrm{fine}}^2 I_r = I_r S_{i,\mathrm{coarse}}.$$

If the mass substeps also intertwine, the identities compose through the entire execution order $M - X - Y - Z - Y - X - M$. The refined walk here replaces each coarse spatial factor by two half factors and each coarse local factor by its supplied subdivision. The result is exact on the framed subspace.

The construction's internal cube links are pure frame ratios and therefore locally flat. Nevertheless, when the coarse x and y translations commute, the large two-by-two fine Wilson loop equals the coarse plaquette. Coarse curvature survives the refinement. Hence exact straight refinement can preserve a nontrivial coarse plaquette; it does not require that every loop become flat.

12.2 A flat cube need not reconstruct the prescribed edge

Conversely, take an arbitrary ordered half-link pair (a, b) and call its actual product $v = ba$. It is always a subdivision of v, even when the separately prescribed coarse link is $u \neq v$. The exact directional defect is

$$\begin{aligned} D_{\text{fine}}\psi &:= S^2_{A,\sigma,(a,b)} I_a \psi - I_a S_{A,\tau,u}\psi \\ &= I_a (S_{A,\tau,v} - S_{A,\tau,u})\psi. \end{aligned} \tag{57}$$

The equality follows by applying the exact theorem to the product v and using linearity of I_a. Endpoint restriction makes I_a injective, so the fine defect vanishes on a state precisely when its coarse mismatch vanishes on that state. This statewise statement is not a claim that one test state detects every link error.

A concrete counterexample removes any possible inference from flatness alone. Let both half-links be i, and prescribe $u = 1$. Their product is $v = -1$. A cube whose twelve positive links are all i has every face plaquette equal to one. On a one-site periodic lattice, however, $\tau = 1$ and

$$S_{A,1,1} = P_+ + P_- = 1, \qquad S_{A,1,-1} = -P_+ - P_- = -1.$$

Thus the coarse mismatch is -2ψ, and the normalized fine defect has norm $2\,\|\psi\|$. The same link data are perfectly flat in the path-comparison sense and maximally wrong relative to the prescribed straight link. This is the role of the quarter-turn witness in the formal source.

Two questions must therefore remain separate: do the ordered fine links reconstruct the desired coarse edge, and do alternative routes between the same endpoints agree? The first is an adjacent-scale compatibility question; the second measures curvature. A reliable error estimate keeps both contributions visible.

13 Hilbert estimates with explicit constants

Finite-dimensional continuity is not enough to identify an operator norm with physical ℓ^2 energy. The earlier continuous-linear packaging retains the ambient norm on a function space. The Hilbert scalar layer instead realizes a finite four-spinor field as

$$\mathcal{H}_S = \ell^2(S \times \{0,1,2,3\}; \mathbb{C}), \qquad \|\psi\|^2 = \sum_{x \in S} \|\psi(x)\|^2_{\mathbb{C}^4}.$$

The formal realization is `EuclideanSpace`; its norm is identified explicitly with the real part of the previously defined state inner product. All operator norms from this point use this Hilbert structure. Assume S is finite and nonempty when taking a maximum over sites.

Theorem 13.1 (Constant-one directional mismatch bound)**.** *Let A be a Hermitian Clifford involution, τ a site bijection, and $u, v : S \to U(1)$. Then*

$$\left\|S_{A,\tau,v} - S_{A,\tau,u}\right\| \le \delta(u,v) := \max_{x\in S} |v(x) - u(x)| = \max_{x\in S} |v(x)u(x)^{-1} - 1|. \tag{58}$$

The constant is independent of the number of sites. The actual half-grid defect in (57), with half-cell Hilbert measure, obeys the same bound.

Proof. At each output site the difference of the shifts has a positive-channel term and a negative-channel term. Orthogonality gives the exact energy identity

$$\begin{aligned}\|(S_v - S_u)\psi\|^2 = \sum_x [&|v(\tau^{-1}x) - u(\tau^{-1}x)|^2 \left\|P_+\psi(\tau^{-1}x)\right\|^2 \\ &+ |v(x)^{-1} - u(x)^{-1}|^2 \left\|P_-\psi(\tau x)\right\|^2].\end{aligned}$$

There is no mixed term, even though the two spinors come from different sites: they lie in orthogonal fixed channel ranges. Unit modulus implies $|v^{-1} - u^{-1}| = |v - u|$. Bound every coefficient by δ^2, reindex the two sums through τ, and use $\|P_+\eta\|^2 + \|P_-\eta\|^2 = \|\eta\|^2$. The result is at most $\delta^2 \|\psi\|^2$. Taking square roots and then the operator supremum proves the first inequality. Multiplication by $u(x)^{-1}$ proves the equality of scalar maxima. Finally (57) and the isometry of I_a identify the fine defect norm on every input with the coarse mismatch norm. □

In ordinary Euclidean coordinates on the fine site set, the half-cell measure is implemented by multiplying the interpolated coordinates by $1/\sqrt{2}$. Omitting this factor changes the norm and inserts an artificial $\sqrt{2}$ into the comparison. The quarter-turn example attains the value two, showing that the constant one cannot be reduced uniformly over all allowed data.

13.1 The telescoping mechanism for a complete walk

Let $I : X \to Y$ be a bounded linear comparison map and let $F_j : Y \to Y$, $C_j : X \to X$ be fine and coarse factors. Write $D_j = F_j I - I C_j$. Two factors satisfy the exact identity

$$F_2F_1I - IC_2C_1 = F_2D_1 + D_2C_1.$$

Repeated insertion gives, in execution order $1, \dots, N$,

$$F_N \cdots F_1 I - I C_N \cdots C_1 = \sum_{j=1}^{N} F_N \cdots F_{j+1} D_j C_{j-1} \cdots C_1. \tag{59}$$

Empty products are identities. If every F_j, C_j is contractive, the triangle inequality and submultiplicativity yield

$$\|F_N \cdots F_1 I - I C_N \cdots C_1\| \le \sum_{j=1}^{N} \|D_j\|.$$

No separate bound on I is needed for this form, since its contribution is already present inside each D_j. Exact local intertwining makes every term zero.

For the scalar seven-factor estimate, both walks act on the same finite Hilbert space, with $I = 1$. One uses the reconstructed links v_x, v_y, v_z, the other the prescribed links u_x, u_y, u_z, and both use the same

two contractive mass factors M_0, M_1. With execution order $M_0 - X - Y - Z - Y - X - M_1$, define $\delta_i = \delta(u_i, v_i)$. The mass defects vanish. Equation (59) and (58) prove

$$\|W_v - W_u\| \le 2\delta_x + 2\delta_y + \delta_z. \tag{60}$$

The repeated X and Y factors explain the coefficients two. The upper bound need not be attained by the complete walk: its individual error terms can cancel. It is an operator estimate between effective common-space walks. The separate half-grid theorem supplies the genuine normalized fine directional estimate; a general spatially varying three-axis comparison still has to provide its actual embeddings and local factor defects.

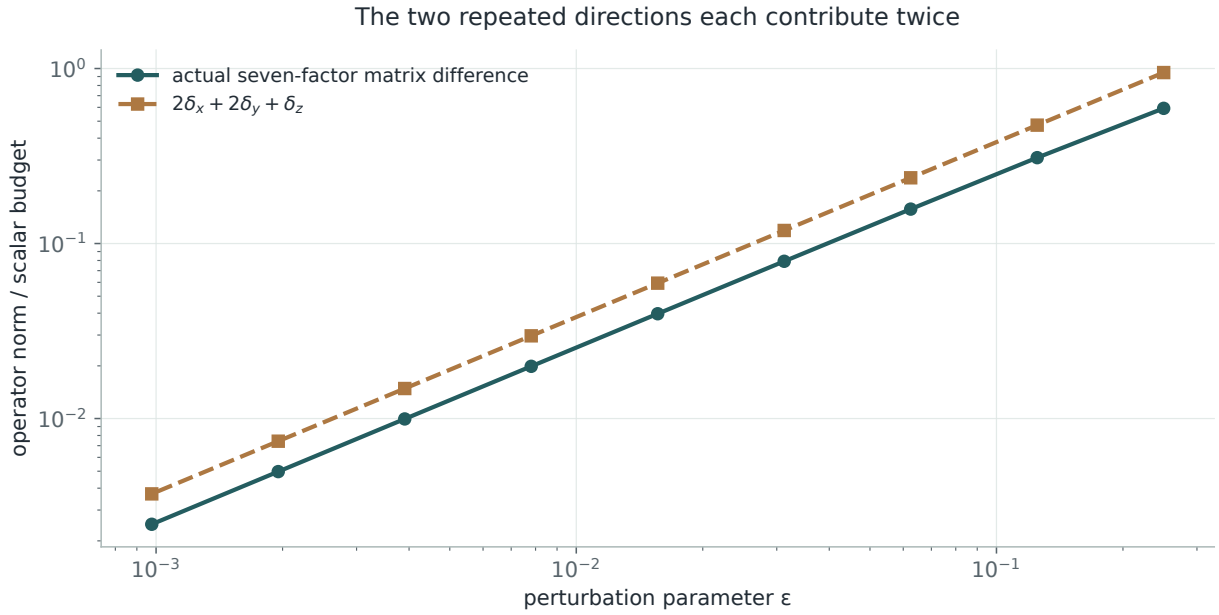


Figure 9. A finite matrix illustration of the telescoping mechanism and the repeated-factor budget. The plotted common-space example tests the algebraic estimate; it is not an additional formal lattice-to-continuum theorem.

14 A two-stage edge and path estimate

The cube phase discrepancies can be made into bounded operators on the same finite Hilbert state space. For $z \in U(1)$, let $P_z\psi = z\psi$ uniformly at every site. If $p = \Omega q$, then

$$P_p - P_q = q(\Omega - 1)\mathrm{id}, \qquad \|P_p - P_q\| = |\Omega - 1|.$$

The norm equality uses a nonzero Hilbert space. It is the operator form of (54).

The formal construction selects five ordered phase pairs, in this order:

$$(p_{XYZ}, p_{YXZ}), \quad (p_{YXZ}, p_{YZX}), \quad (p_{YZX}, p_{ZYX}), \quad (p_{XYZ}, p_{XZY}), \quad (p_{XZY}, p_{ZXY}).$$

Let Q_f be the execution-order product of the first phase operators in these pairs, and Q_c the product of the second ones. Identity factors can be placed at the ends to fit the seven-factor interface. Applying (59) yields

$$\|Q_f - Q_c\| \le \kappa(\ell).$$

These are explicitly constructed products of uniform phase operators. Calling them a path-reordering stage describes their algebraic role; it does not assert that a global spatially varying inter-cell assembly has already been constructed.

Theorem 14.1 (Combined finite Hilbert estimate)**.** *With the common-space walks and contractive mass factors from* (60)*, define*

$$D_{\text{total}} = Q_f W_v - Q_c W_u.$$

Then

$$\|D_{\text{total}}\| \le 2\delta_x + 2\delta_y + \delta_z + \kappa(\ell). \tag{61}$$

If $v_i = u_i$ for all three directions and the cube is flat, the defect is exactly zero.

Proof. Insert and subtract $Q_f W_u$:

$$D_{\text{total}} = Q_f(W_v - W_u) + (Q_f - Q_c)W_u.$$

All phase factors are unitary, and W_u is a product of contractions. Taking norms bounds the first term by the straight-edge budget and the second by the path budget. In the exact sector the walks agree and every selected phase pair agrees, so both terms vanish before taking norms. □

The sum keeps the logical independence proved in Section 12. A bound on face curvature cannot replace a bound on reconstruction ratios. Conversely, exact ordered products do not remove a nonflat alternative-path discrepancy. Neither conclusion requires that the estimate be sharp for every configuration.

15 Connection-generated links and a quadratic mesh bound

The scalar budgets become useful when links arise from a specified background rather than independent phases. Consider real abelian connection components A_x, A_y, A_z on a Euclidean chart. Define exact coordinate-edge angles and forward links by

$$\theta_i(p,h) = \int_0^h A_i(p + te_i)\,dt, \qquad u_i(p,h) = e^{i\theta_i(p,h)}. \tag{62}$$

The coupling is absorbed into A_i. These links are automatically unitary, and the interval integral is oriented when $h < 0$.

15.1 Exact composition precedes every estimate

Continuity of A_i along its own coordinate lines permits interval-integral additivity and a change of variable:

$$\begin{aligned}\theta_i(p,h) &= \int_0^{h/2} A_i(p+te_i)\,dt + \int_{h/2}^{h} A_i(p+te_i)\,dt \\ &= \theta_i(p,h/2) + \theta_i(p+(h/2)e_i, h/2).\end{aligned}$$

Exponentiation gives the ordered identity

$$u_i(p+(h/2)e_i, h/2)u_i(p,h/2) = u_i(p,h). \tag{63}$$

Thus the reconstructed and prescribed links agree exactly and all three straight-edge errors vanish. This is an additivity theorem for exact line integrals. Independently approximating each integral by a quadrature rule would introduce another error term, which is not silently included in the present bound.

For an affine component, $A_i(p+te_i) = A_i(p) + t\,\partial_i A_i$, so

$$\theta_i(p,h) = hA_i(p) + \tfrac{1}{2}h^2\partial_i A_i.$$

The face angle is exactly h^2F_{ij}, where $F_{ij} = \partial_i A_j - \partial_j A_i$. Hence $\Omega_{ij} = e^{ih^2F_{ij}}$. Using $|e^{it}-1| \le |t|$ and the selected five faces yields

$$\|D_{\text{total}}(h)\| \le h^2(|F_{xy}| + 2|F_{xz}| + 2|F_{yz}|). \tag{64}$$

This exact affine calculation is stronger than a truncated Taylor argument. It also fixes the orientation sign of the curvature convention.

15.2 The actual coordinate regularity assumptions

The more general formal connection structure records the three components, the six cross derivatives $\partial_i A_j$ for $i \neq j$, and their one-dimensional derivative identities along the relevant coordinate lines. Each component is continuous along its own integration direction. Each cross derivative is continuous along the line in its differentiation direction. Finally there are global nonnegative constants

$$|\partial_i A_j(p)| \le B_{j|i} \qquad (i \neq j).$$

For example, $B_{y|x}$ bounds the derivative of A_y in the x direction. These explicit hypotheses are sufficient for the coordinate arguments below. They avoid appealing to an unspecified stronger multivariable smoothness class.

The fundamental theorem of calculus first gives

$$|A_j(p+he_i) - A_j(p)| \le B_{j|i}|h|.$$

Integrating along the transverse j edge then gives

$$|\theta_j(p+he_i,h) - \theta_j(p,h)| \le B_{j|i}|h|^2. \tag{65}$$

Indeed the integrand difference is bounded by $B_{j|i}|h|$ at every point of an interval of length $|h|$. Oriented integration changes the sign but not this absolute bound.

For the xy face, the two-path angle difference is

$$\Theta_{xy}(p,h) = \theta_y(p+he_x,h) - \theta_y(p,h) \\ - (\theta_x(p+he_y,h) - \theta_x(p,h)).$$

Therefore

$$|\Theta_{xy}(p,h)| \le (B_{y|x} + B_{x|y})h^2, \qquad |\Omega_{xy}(p,h) - 1| \le (B_{y|x} + B_{x|y})h^2.$$

For the last implication, $e^{it} - 1 = \int_0^t ie^{is}\,ds$ has absolute value at most $|t|$. The other two plane estimates follow from their corresponding cross derivatives. This proof bounds absolute derivatives separately; it does not prove a sharper general surface-integral formula that exploits cancellation inside F_{ij}.

Theorem 15.1 (Regular background estimate)**.** *Suppose the coordinate regularity and derivative bounds above hold. Set*

$$\mathcal{A}_\partial = B_{y|x} + B_{x|y} + 2(B_{z|x} + B_{x|z}) \\ + 2(B_{z|y} + B_{y|z}). \tag{66}$$

Generate the reconstructed and prescribed finite-site links from the exact integrals (62)*, and generate the twelve links of the comparison cube from the same connection at a chart point p. For the two-stage finite Hilbert operator of Section 14, with the same contractive mass factors on both sides,*

$$\|D_{\text{total}}(h)\| \le h^2 \mathcal{A}_\partial. \tag{67}$$

Proof. Equation (63) makes $\delta_x = \delta_y = \delta_z = 0$. Apply the appropriate face estimate to each term in (55). There is one xy term, two xz terms and two yz terms, giving $\kappa \le h^2\mathcal{A}_\partial$. Substitute in (61). □

The formal theorem permits an arbitrary finite nonempty site type, three site bijections and a chart-position map, along with an independently specified cube basepoint. It does not impose a geometric identity between every site translation and that position map. The result is consequently an exact estimate for the declared generated operators; interpreting them as one global geometric lattice requires those additional compatibility choices.

15.3 A spatially varying witness

Take $A_x = A_z = 0$ and $A_y(x,y,z) = \sin x$. Its curvature $F_{xy} = \cos x$ is nonconstant. The only nonzero derivative envelope is $B_{y|x} = 1$, so $\mathcal{A}_\partial = 1$. Here

$$\theta_y(p,h) = h\sin x, \qquad \Omega_{xy}(p,h) = \exp(ih[\sin(x+h) - \sin x]).$$

The other plane holonomies are one, and the two-stage defect is bounded by h^2. The bound is uniform in the basepoint. The trigonometric formula illustrates how actual phases can be substantially smaller than a global envelope; it does not change the envelope's role in a uniform theorem.

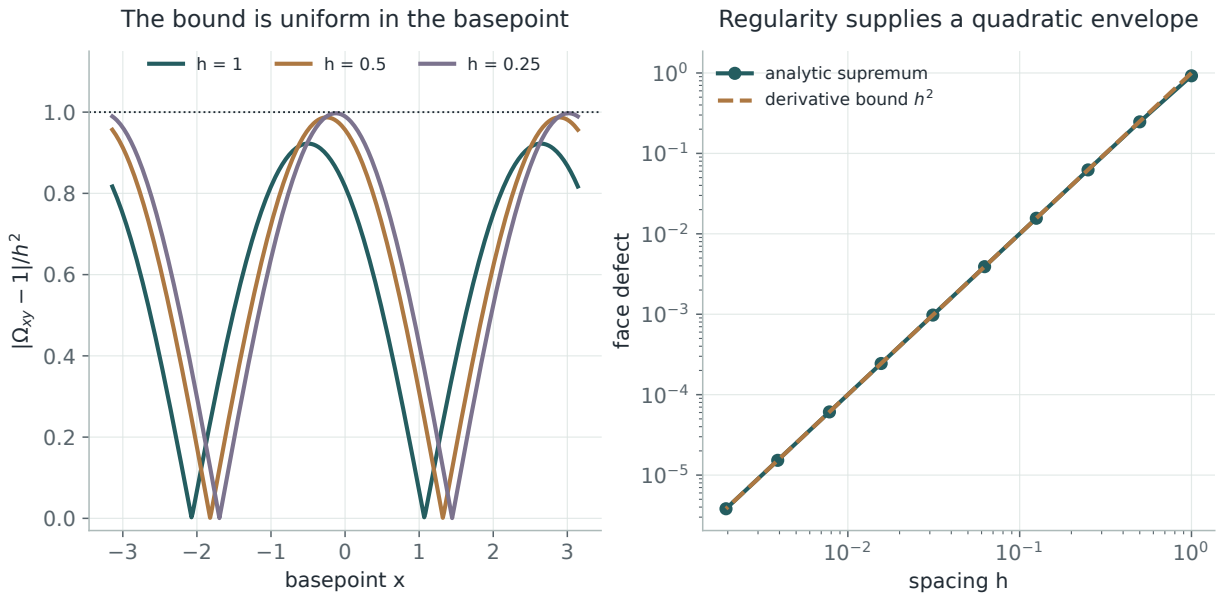


Figure 10. The non-affine background $A_y = \sin x$ has varying curvature and exact line-integral links. The quadratic envelope follows from $|\partial_x A_y| \le 1$; the computed phase amplitudes illustrate this bound over the plotted finite range.

16 Summable refinement and what a limit still requires

For the dyadic mesh $h_n = a_0 2^{-n}$, the regular estimate produces the nonnegative sequence

$$\varepsilon_n = C4^{-n}, \qquad C = a_0^2 \mathcal{A}_\partial.$$

The actual generated operator norms are bounded by this sequence and hence are summable. This is more information than merely $\varepsilon_n \to 0$. A sequence with adjacent errors tending to zero need not converge; summability controls the accumulated tail.

Proposition 16.1 (A complete-space limit with an explicit tail)**.** *Let f_n be a sequence in a complete metric space and suppose*

$$d(f_n, f_{n+1}) \le C4^{-n} \qquad (n \ge 0), \quad C \ge 0.$$

Then f_n converges to some f_∞, and

$$d(f_n, f_\infty) \le \sum_{k=n}^{\infty} C4^{-k} = \tfrac{4}{3} C4^{-n}. \tag{68}$$

Proof. For $m > n$, the triangle inequality gives

$$d(f_n, f_m) \le \sum_{k=n}^{m-1} C4^{-k} \le \tfrac{4}{3} C4^{-n}.$$

The upper bound tends to zero with n, so the sequence is Cauchy. Completeness provides a limit. Passing $m \to \infty$ in the distance estimate proves the tail bound, and the displayed evaluation is the geometric series. □

The formal regular-connection limit theorem has precisely this conditional shape: it takes a sequence in a named complete metric space and an adjacent-distance hypothesis bounded by the regular dyadic budget. The source also proves summability of the actual connection-generated finite-operator norms. These are two compatible ingredients, but the operator bound alone does not manufacture the sequence f_n or identify its terms across different Hilbert spaces.

For a physical refinement application one must define the reconstructed states or evolutions in a common space and derive the required distance estimate from the relevant operator defects. If the input state is not normalized, its norm enters that derivation. If the number of time steps grows with the refinement level, the accumulation in time must also be estimated. A limit supplied by this abstract argument is not automatically a Dirac propagator or a solution of a coupled gauge-field system.

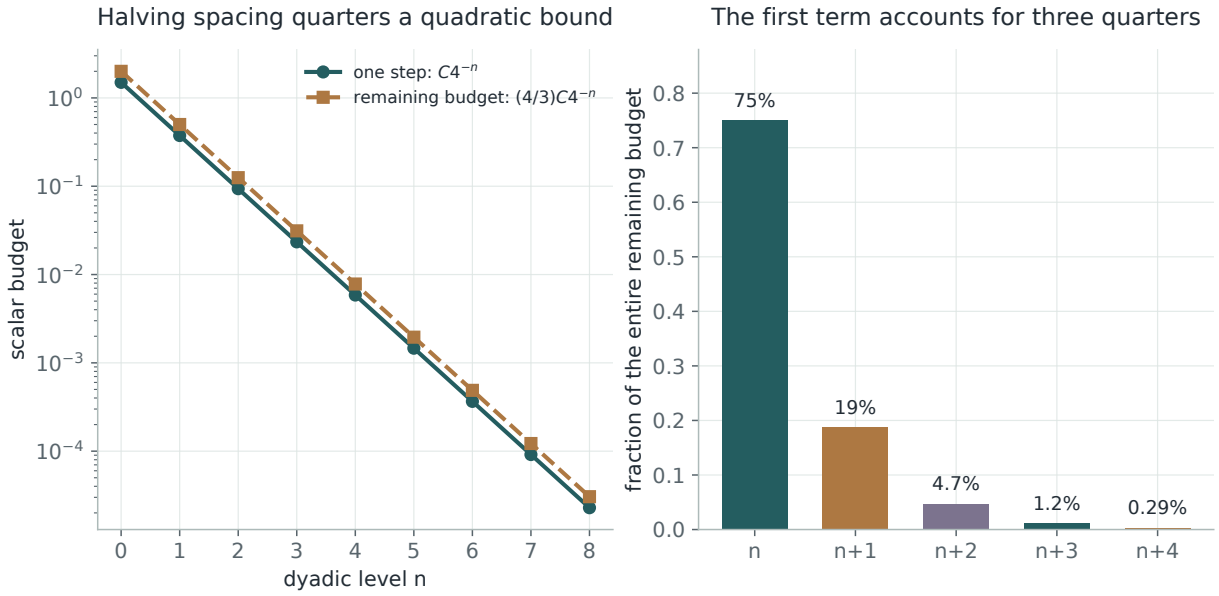


Figure 11. For the regular connection envelope, halving the mesh quarters the one-level budget. The remaining accumulated error is bounded by the geometric tail $\frac{4}{3}C4^{-n}$ once the adjacent-distance hypothesis has been established in a common complete space.

17 From weighted cells to a continuum Hilbert space

The reconstruction blocks that follow the connection estimates supply an important part of the common-space infrastructure. They establish what finite cell coefficients mean as L^2 fields, then identify a concrete family of approximating subspaces. They do not by themselves transfer the gauge-coupled walk or its generator through those maps.

17.1 Cell values and normalized coefficients

For a measurable finite partition into cells of equal volume $v > 0$, distinguish the physical constant value $q_a \in \mathbb{C}^4$ on cell C_a from the normalized coefficient $\psi_a = \sqrt{v}\, q_a$. The step interpolation is

$$(J\psi)(x) = v^{-1/2}\psi_a \quad (x \in C_a),$$

with equality understood almost everywhere. Direct integration gives

$$\|J\psi\|_{L^2}^2 = \sum_a \int_{C_a} v^{-1} \|\psi_a\|^2 = \sum_a \|\psi_a\|^2 .$$

Equivalently, physical cell values carry the weighted norm $v \sum_a \|q_a\|^2$. The adjoint averaging map yields the orthogonal projection onto the finite step subspace. Its Pythagorean decomposition proves best approximation in L^2.

At octree depth n, the number of cells is 8^n, their side scale is 2^{-n} and their volume is 8^{-n}. The first construction uses infinite octant addresses with the uniform product probability measure. Finite-prefix step subspaces are nested and have dense union. The weighted normalization is thus tied to the actual measure of a prefix cell, rather than inferred from the number of coordinates alone. The later source convention `euclideanDyadicLevel d` uses prefix depth $d+1$; the geometric depth n used here counts digits directly.

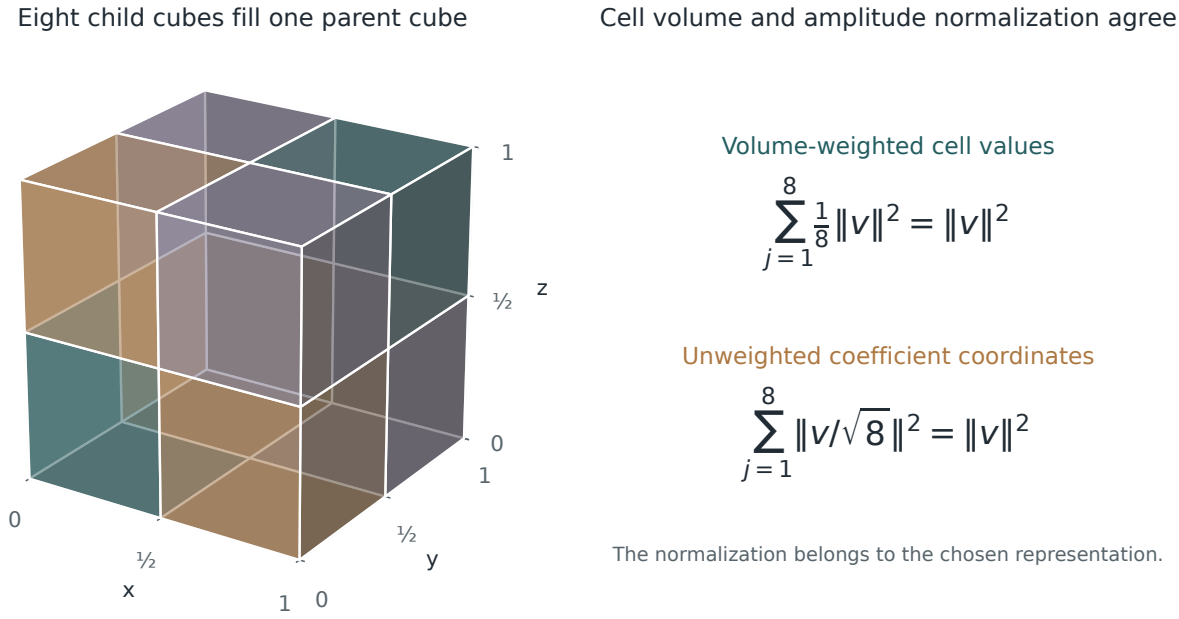


Figure 12. Eight equal children preserve a parent cell's total volume. Constant transported values use child volume weights; counting-normalized coefficients instead include the square root of the child volume. These are equivalent descriptions of the same Hilbert normalization.

17.2 The geometric identification is an additional theorem

Splitting each octant digit into three bits and applying positional binary decoding gives a continuous surjection from address space onto the unit cube. The decoder is not pointwise injective: dyadic coordinates admit the familiar terminating and nonterminating binary descriptions. The measure result then proves that fair binary product measure pushes forward to cube volume, producing a norm-preserving L^2 pullback.

At the intermediate stage, finite-resolution Euclidean subspaces are defined by adjoint transport of address-space levels. Density and strong convergence of their orthogonal projections can be proved at that stage, but it would be premature to call them ordinary dyadic-box subspaces. The next block closes

this identification: a prefix cell equals the preimage of its corresponding concrete Cartesian dyadic box almost everywhere. The formal boxes are closed and share boundary faces; those faces have zero volume. Equality in L^2 therefore survives the binary boundary ambiguity.

The transported and concrete closed levels are then proved equal. Their union is dense, and the orthogonal projections P_n satisfy

$$\|P_n f - f\|_{L^2} \longrightarrow 0 \qquad \text{for every } f \in L^2([0,1]^3; \mathbb{C}^4).$$

This is a state-space approximation theorem. Strong convergence of projections does not control an unbounded Dirac operator, establish graph convergence, or identify the limit of gauge-coupled propagators. Those statements need consistency on suitable domains and stability estimates for the transported operators.

The resulting position is precise. Exact covariant transport, genuine geometric subdivisions, finite Hilbert defect bounds, connection-derived mesh estimates and concrete continuum approximation spaces are available as separate, source-backed layers. Their assembly into a regular-background physical Hamiltonian remains a further mathematical task. The self-adjoint free-torus Hamiltonian provides an operator endpoint for the ungauged branch; importing these background estimates into that setting requires explicit domain and comparison-map work.

18 Gauge transport and its continuum comparison

Endpoint cancellation gives covariance of composed transport and conjugation of closed holonomy. In the specified finite Maxwell–Dirac system, matter continuity, Gauss propagation, and electric work fit together exactly. Ordered non-Abelian products retain the path information required by refinement.

Geometric interpolation supplies actual midpoints and weighted isometries. Straight-edge reconstruction and path curvature enter distinct error budgets for explicitly constructed finite operators. Regular line-integral links reconstruct straight edges exactly, while derivative envelopes give summable dyadic bounds for the chosen curvature defect.

Correctly normalized finite fields become L^2 step functions. Concrete dyadic boxes identify the approximating subspaces modulo null boundaries, and their projections converge strongly on each fixed field. Comparison of physical generators and evolution remains a separate task.

18.1 The remaining analytic bridge

The next background-coupled result needs a common operator setting. One must specify the physical connection, its regularity, the continuum spinor domain, the discrete operators, and the embeddings that compare them. Consistency on a controlled core, stability, and a suitable graph or resolvent argument then have to connect the finite refinement estimates to the continuum operator. Each limit must identify its topology and its dependence on the state or its regularity.

The free Dirac route already provides a useful endpoint for this comparison. The free construction gives a self-adjoint physical torus Hamiltonian on its weighted Fourier domain. The generator chapter supplies its spectral-evolution identification as a written proof. Its formal certification remains distinct from the recorded operator certificate. A background gauge extension must account for those operator-domain and evolution questions as well as the link estimates developed here.

18.2 From a prescribed connection to coupled dynamics

A background estimate controls a supplied field. A coupled gauge theory makes that field part of the unknown. The finite semidiscrete example illustrates the desired compatibility among matter transport, electric evolution, and a Gauss constraint. Extending this compatibility to a specified continuum setting requires an existence theory, constraint propagation, and a convergence statement for the coupled system. Local covariance alone supplies none of these additional conclusions.

The constrained-gravity chapter examines the additional obligations imposed by a dynamical geometry: regular-chart reduction, constraint closure, boundary terms, and strong-field evolution. Arithmetic refinement and cosmological transport retain the same requirement to identify their state spaces and comparison maps, without sharing all of the gauge model's assumptions.

> **The next theorem should name its comparison**
>
> Specify the two state spaces, the operators or evolutions to be compared, the map between resolutions, and the norm or topology in which the defect tends to zero. Then state which connection and domain hypotheses make that comparison valid.

Sources and further reading

Project sources

[P1] Komorebi P. *Projector Refinement Physics, Volume I*. The Abelian gauge and ordered non-Abelian holonomy studies supply the finite field equations and matrix identities. These are written mathematical arguments. Exact formal coverage is identified separately.

[P2] Komorebi P. *Komorebi Physics formal source*. The gauge verification appendix records the transport and refinement declarations, their hypotheses, and the limitations of the retained compiler evidence.

[P3] Komorebi P. *Projectors, Packets, and Transfer* and *Fourier Packets and the Continuum Dirac Operator*. The packet framework and free-field operator construction in this book. Extending the latter to a background connection requires the domain and evolution comparisons stated here.

Primary context

[C1] C. N. Yang and R. L. Mills. "Conservation of Isotopic Spin and Isotopic Gauge Invariance." *Physical Review* **96** (1954), 191–195. https://doi.org/10.1103/PhysRev.96.191.

[C2] Kenneth G. Wilson. "Confinement of quarks." *Physical Review D* **10** (1974), 2445. https://doi.org/10.1103/PhysRevD.10.2445.

[C3] John Kogut and Leonard Susskind. "Hamiltonian formulation of Wilson's lattice gauge theories." *Physical Review D* **11** (1975), 395. https://doi.org/10.1103/PhysRevD.11.395.

[C4] Pablo Arnault and Fabrice Debbasch. "Quantum walks and discrete gauge theories." *Physical Review A* **93** (2016), 052301. https://doi.org/10.1103/PhysRevA.93.052301.

[C5] C. Cedzich, T. Geib, A. H. Werner, and R. F. Werner. "Quantum walks in external gauge fields." *Journal of Mathematical Physics* **60** (2019). Author preprint arXiv:1808.10850v1, 31 August 2018. https://arxiv.org/abs/1808.10850v1; https://doi.org/10.1063/1.5054894.

How these sources are used

Yang and Mills provide the original continuum setting for local non-Abelian covariance. Wilson and Kogut–Susskind identify the lattice and Hamiltonian settings in which links, closed-loop observables, electric variables, and constraints belong. These references provide context; their quantization or confinement results are not premises of the elementary identities proved here. A gauge-invariant expression for a walk does not by itself supply a Hamiltonian phase space for its gauge field.

Arnault–Debbasch and Cedzich–Geib–Werner–Werner are prior work on gauge-covariant discrete-time walks. In the latter preprint, equations (19)–(20) give endpoint covariance and conjugation of closed transport, while Section IV.6 explains why coupling depends on the ordered decomposition of a walk. These are direct points of comparison with the conventions and factorwise refinement used here. Their plaquette in equation (22) traverses the opposite orientation to the ratio of forward paths used in this volume; in the Abelian case the two phases are inverses. Its global lattice statements should not be transferred to a periodic quotient without accounting for noncontractible loop holonomy.

The source-check record separates inspected full text from bibliographic authentication. Relevant full-text sections of the Yang–Mills article and the Cedzich et al. preprint were inspected. For Wilson, Kogut–Susskind, and Arnault–Debbasch, the accompanying record authenticated the publisher metadata and abstract, but full-text access through the publisher was unavailable. No uninspected result is used to discharge a proof obligation. This is a focused source comparison, not an exhaustive literature review or a priority claim.

Constrained Gravity and Strong Fields

Constraint algebra, regular reduction, and spherical geometry

Geometry, constraints, and strong fields

Gauge covariance identifies compatible transport. Gravity asks for more: the geometry, constraints, and evolution must be specified together. A propagation rule, a canonical constraint algebra, and an evolved matter system are different mathematical objects, with different hypotheses.

This account proceeds from weighted transport geometry and reduced Gowdy evolution to finite linearized ADM constraints. The nonlinear obstruction explains why a nonzero derivative cannot satisfy the ordinary Leibniz rule on a complete finite pointwise algebra. Matrix moment maps, homological transfer, and spectral representations provide alternatives in their stated settings.

A real Lorentzian action supplies a classical gauge-theoretic parent. Boundary constructions and local second-class reduction retain their regularity and nondegeneracy conditions. The spherical sector makes mass, flux, and horizon diagnostics explicit, allowing exact identities to be compared with bounded numerical experiments.

The results do not amount to a complete continuum or quantum theory of gravity. A manufactured electrovac trajectory is not an existence theorem for nonlinear collapse. A scalar trapped-sphere diagnostic is not a continuum formation theorem without stability, convergence, and event control. These distinctions are part of the mathematics rather than qualifications added after it.

The shared resolution language continues in the arithmetic and polyphase constructions that follow. It makes the state spaces, comparison maps, and residual obstructions explicit without equating their physical interpretations.

Reading routes and conventions

For the geometric route, begin with transport and local frames, then follow the Gowdy model and the finite linearized ADM system. The nonlinear obstruction explains why a product rule that is harmless in a continuum calculation becomes a substantive issue on a finite pointwise algebra.

For the constrained-system route, begin with the canonical bracket and the exact finite moment map. Continue through homological transfer, resolved spectral evaluation, the classical BV parent, and the boundary and local reduction conditions. A field-dependent change of ghost variables and a canonical transformation of the extended phase space must be distinguished. The latter carries the appropriate momentum terms with it.

For the strong field route, begin with the spherical mass identity. The charged horizon calculations and the scalar mass-flux compatibility give concrete statements that can be checked independently. The numerical records then test selected trajectories and initial data. A manufactured gauge trajectory of a fixed spacetime and an evolved scalar-collapse model answer different questions.

Convention	Meaning
Canonical brackets	Each sector declares its variables and sign. In standard coordinates, $\{F, G\} = \sum_a (F_{q_a} G_{p_a} - F_{p_a} G_{q_a})$, so $\dot{F} = \{F, H\}$.
Geometry and signature	Lorentzian signature, curvature orientation, density weights, lapse, shift, and tetrad conventions are stated where they enter. Local transport alone does not select an Einstein action.
Gauge and reduction	Gauge choice, constraint enforcement, and removal of second-class variables are separate operations. An invertible gauge-fixing or Dirac matrix defines the regular chart used in the corresponding argument.
Spectral resolution	K denotes retained degree and N the sampling resolution in the resolved constructions. Conditions such as $16K < N$ are sufficient degree budgets in a declared representation. They are not invariant under every nonlinear change of variables.
Locality	A finite-range stencil, a full spectral evaluation, a matrix commutator, and a radial prefix constraint solve have different dependence on the data. The implementation is named before its locality is assessed.
Time stepping	The conservation proof for a projected discrete-gradient step concerns an autonomous Hamiltonian and an exactly solved implicit step. Approximate solves retain a residual that must be measured.
Spherical units	Geometrized units are used in the spherical examples. The mass, charge, radial coordinate, and slicing branch are defined before horizon or flux formulas are used.
Evidence	Ordinary proofs, inherited executable records, newly run bounded computations, and archived formal declarations retain their own scope. Numerical plots do not substitute for an existence or convergence proof.

Geometry and constrained evolution

1 Geometric transport and local work

The first gravitational question is local: how does a prescribed geometry change transport and energy exchange? A second question is dynamical: which equations determine that geometry? Keeping these questions distinct lets a finite model establish an exact balance law before it is coupled to a gravitational constraint system. Throughout this section the metric coefficients are supplied functions. The scalar field evolves on them.

1.1 A finite scalar action on a diagonal metric

Take a periodic chain with $N \geq 3$ sites, spacing $h > 0$, and indices modulo N. The continuum metric motivating its coefficients is

$$ds^2 = -N(t,x)^2 dt^2 + a(t,x)^2 dx^2, \qquad N(t,x) > 0, \quad a(t,x) > 0.$$

The repeated letter N in this metric is the lapse, whereas the integer N counts lattice sites. In the finite formulas below only the coefficient names w, c, and v are needed:

$$w_j = \frac{a_j}{N_j} > 0, \qquad c_{j+1/2} = \frac{N_{j+1/2}}{a_{j+1/2}} > 0, \qquad v_j = N_j a_j m^2 \geq 0.$$

Thus N/a is the coordinate characteristic speed of the continuum massless equation. Assume these finite coefficient arrays are continuously differentiable in time. For a real field $\varphi_j(t)$ set $\Delta_j = \varphi_{j+1} - \varphi_j$ and define

$$L_h = \frac{h}{2} \sum_j w_j \dot{\varphi}_j^2 - \frac{1}{2h} \sum_j c_{j+1/2} \Delta_j^2 - \frac{h}{2} \sum_j v_j \varphi_j^2. \tag{1}$$

Variations with compact time support give

$$w_j \ddot{\varphi}_j + \dot{w}_j \dot{\varphi}_j = \frac{c_{j+1/2}\Delta_j - c_{j-1/2}\Delta_{j-1}}{h^2} - v_j \varphi_j. \tag{2}$$

The neighboring edge terms in the action produce the two terms in the spatial numerator. The time integration by parts differentiates $w_j \dot{\varphi}_j$, including the $\dot{w}_j \dot{\varphi}_j$ term.

Proposition 1.1 (Local energy, flux, and geometric work)**.** *For every twice differentiable solution of* (2)*, define*

$$\begin{aligned} e_j &= \frac{h}{2} w_j \dot{\varphi}_j^2 + \frac{h}{2} v_j \varphi_j^2 + \frac{c_{j+1/2}}{4h} \Delta_j^2 + \frac{c_{j-1/2}}{4h} \Delta_{j-1}^2, \\ F_{j+1/2} &= -\frac{c_{j+1/2}}{2h} \Delta_j (\dot{\varphi}_j + \dot{\varphi}_{j+1}), \\ s_j &= -\frac{h}{2} \dot{w}_j \dot{\varphi}_j^2 + \frac{h}{2} \dot{v}_j \varphi_j^2 + \frac{\dot{c}_{j+1/2}}{4h} \Delta_j^2 + \frac{\dot{c}_{j-1/2}}{4h} \Delta_{j-1}^2. \end{aligned} \tag{3}$$

Then

$$\dot{e}_j + F_{j+1/2} - F_{j-1/2} = s_j. \tag{4}$$

In particular, $E_h = \sum_j e_j$ *satisfies* $\dot{E}_h = \sum_j s_j = -\partial_t L_h$*, where the partial derivative holds* φ *and* $\dot{\varphi}$ *fixed.*

Proof. Put $B_j = c_{j+1/2} \Delta_j$. Using the equation of motion, the derivative of the kinetic and mass terms at site j is

$$\frac{\dot{\varphi}_j}{h} (B_j - B_{j-1}) - \frac{h}{2} \dot{w}_j \dot{\varphi}_j^2 + \frac{h}{2} \dot{v}_j \varphi_j^2.$$

The right edge contribution differentiates to

$$\frac{\dot{c}_{j+1/2} \Delta_j^2}{4h} + \frac{B_j}{2h} (\dot{\varphi}_{j+1} - \dot{\varphi}_j),$$

and the left edge has the analogous expression. Adding these terms leaves $B_j(\dot{\varphi}_j + \dot{\varphi}_{j+1})/(2h)$ on the right and its negative counterpart on the left, plus s_j. These are exactly the flux terms in (4). Periodic summation cancels every interface flux. Each edge coefficient derivative appears twice with weight $1/(4h)$, giving $-\partial_t L_h$. □

Every e_j is nonnegative under the stated coefficient assumptions. An interval of cells $r, \dots, s$ obeys the exact coarse balance

$$\frac{d}{dt} \sum_{j=r}^{s} e_j + F_{s+1/2} - F_{r-1/2} = \sum_{j=r}^{s} s_j.$$

Thus coarse energy is obtained by adding fine cell energies and retaining only the exposed boundary fluxes. This is the same cancellation mechanism as projector transfer, applied here to a classical field energy. Its physical units and constitutive coefficients come from the action.

Example 1.2 (Time dependent expansion)**.** Let the lapse equal one, let $a(t) > 0$ be spatially uniform, and let $m = 0$. Then $w = a$, $c = 1/a$, and

$$\dot{E}_h = -\frac{h}{2} \dot{a} \sum_j \dot{\varphi}_j^2 - \frac{\dot{a}}{2ha^2} \sum_j \Delta_j^2.$$

An increasing a transfers energy out of this coordinate-time field Hamiltonian. The supplied background receives no evolution equation from this calculation. A self-consistent matter–geometry model must also evolve its gravitational variables.

For a smooth field sampled at sites and a smooth coefficient sampled at edge midpoints, the spatial expression is second-order consistent with $\partial_x (c \partial_x \varphi)$. Indeed, expanding at x_j gives

$$c(x_j + h/2)[\varphi(x_j + h) - \varphi(x_j)] - c(x_j - h/2)[\varphi(x_j) - \varphi(x_j - h)] = h^2 (c\varphi_x)_x (x_j) + O(h^4).$$

Uniform bounds on the requisite fourth derivatives make the remainder uniform on a compact spatial region. The exact finite balance holds without any Taylor approximation.

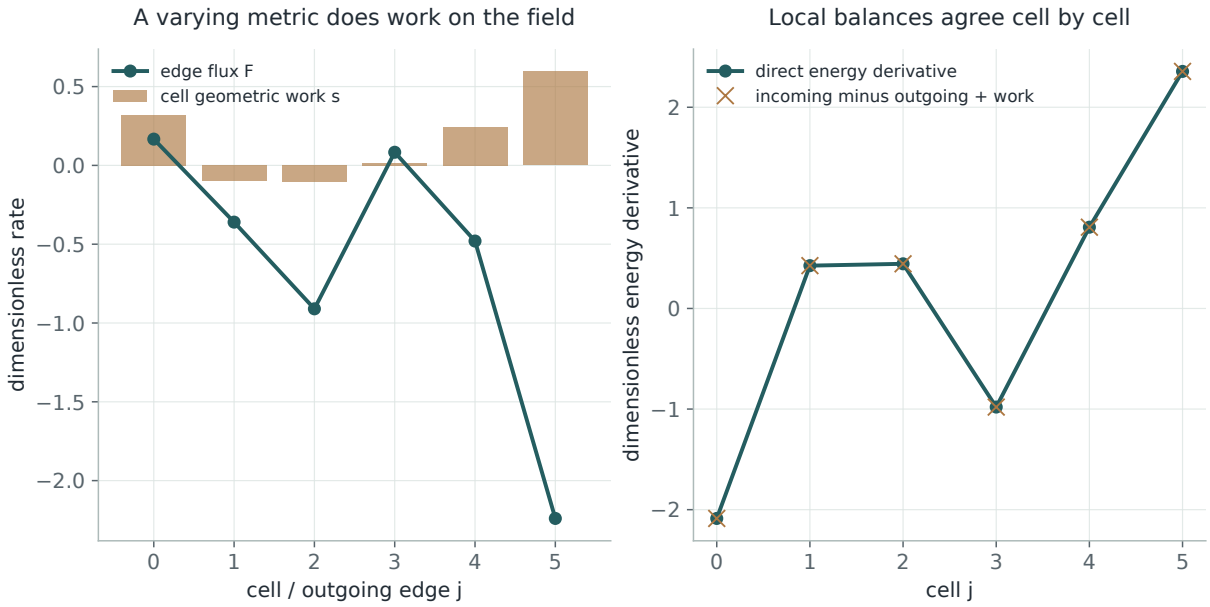


Figure 1. A six-cell example of prescribed metric work. Each direct cell-energy derivative agrees with incoming flux minus outgoing flux plus the coefficient work term. The curves illustrate the exact finite identity for a prescribed background.

1.2 Conformal curvature as an exact boost holonomy

Consider a twice continuously differentiable real function ϕ on an oriented (t, x) rectangle, and the coframe

$$e^0 = e^{\phi} dt, \qquad e^1 = e^{\phi} dx, \qquad ds^2 = e^{2\phi}(-dt^2 + dx^2).$$

Write

$$K = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad \omega = (\phi_x dt + \phi_t dx) K.$$

The coframe satisfies the torsion-free Cartan equations with this connection. Since every connection value is a multiple of the same matrix, $\omega \wedge \omega = 0$. Therefore

$$\Omega = d\omega = (\phi_{tt} - \phi_{xx}) K \, dt \wedge dx, \qquad R = 2e^{-2\phi}(\phi_{tt} - \phi_{xx}). \tag{5}$$

This fixes the scalar-curvature sign convention used in this monograph.

We use the archive's plus-sign edge convention: for each oriented boundary edge define $U_e = \exp(+\int_e \omega)$. This specifies the signed transport observable used below. Multiplying in boundary order yields an exact identity

$$W_{\square} = \exp\left(K \oint_{\partial \square} (\phi_x dt + \phi_t dx) \right) = \exp\left(\frac{K}{2} \int_{\square} e^{2\phi} R \, dt \, dx \right). \tag{6}$$

To see the second equality directly, integrate ϕ_{tt} along the two horizontal endpoints and ϕ_{xx} along the two vertical endpoints, with the boundary orientation from $dt\wedge dx$. This is the rectangular fundamental theorem of calculus. No path-order correction survives because the edge exponents commute.

If the integrated coefficient is χ, then

$$e^{\chi K} = \cosh\chi\, I + \sinh\chi\, K.$$

This preserves the Lorentz form $\mathrm{diag}(-1,1)$ and has positive eigenvalues e^{χ} and $e^{-\chi}$. Its real principal matrix logarithm is χK. These are Lorentz boosts. The positive Hilbert norm used for unitary color transport in the gauge monograph is a different structure.

Example 1.3 (Constant curvature density)**.** For $\phi(t,x) = b(t^2-x^2)/2$, one has

$$\phi_{tt} - \phi_{xx} = 2b, \qquad R = 4b\, e^{-2\phi}, \qquad W_{\square} = e^{2b\ \mathrm{area}(\square)K}.$$

The curvature density $e^{2\phi}R/2$ is constant although R varies with position. Adding oriented subrectangle integrals reconstructs the coarse rapidity exactly.

The plus-sign observable is assigned the endpoint action

$$G(t,x) = e^{\chi(t,x)K}, \qquad U'_{a\to b} = G(b)U_{a\to b}G(a)^{-1}.$$

It corresponds to the signed connection rule $\omega'_{+} = \omega_{+} + d\chi\, K$. Intermediate endpoint factors cancel, so the loop observable is conjugated at its base point. For the Cartan convention $D = d + \omega$, a coframe transformation $e \mapsto Ge$ instead gives $\omega' = \omega - d\chi\, K$ and parallel transport $\exp(-\int_e \omega)$. Its closed-loop rapidity is the negative of the rapidity of $W_{\square}$ defined here. The fixed-frame curvature identity (6) is unchanged. Declaring these signs avoids mixing the archive's edge observable with a different Cartan parallel-transport convention.

2 Weak gravity as a sign and consistency benchmark

Weak gravity provides checks on attraction, variational signs, and optical observables. These checks are performed within the weak-field model. They do not furnish nonlinear gravitational evolution.

2.1 Radial Poisson equation and a finite variational problem

Let $G > 0$, let $\rho(r)$ be a continuous spherical density, and impose regularity $\lim_{r\downarrow 0} r^2\Phi'(r) = 0$. The radial Poisson equation gives

$$(r^2\Phi')' = 4\pi G r^2\rho, \qquad M_{\mathrm{enc}}(r) = 4\pi\int_0^r s^2\rho(s)\,ds, \qquad \Phi'(r) = \frac{GM_{\mathrm{enc}}(r)}{r^2}.$$

The acceleration is $a_r = -\Phi'$. A positive density therefore gives an inward radial acceleration.

Putting $u(r) = r\Phi(r)$ removes the radial first derivative:

$$r^2\Phi' = ru' - u, \qquad (r^2\Phi')' = ru'', \qquad u'' = 4\pi G r\rho.$$

For fixed endpoint values, this equation is the stationary condition of

$$S[u] = \int_0^{R_0} \left[\frac{(u')^2}{2G} + 4\pi r \rho u \right] dr.$$

On $r_i = ih$, $i = 0, \dots, M$, with $Mh = R_0$, its finite counterpart is

$$S_h[u] = \sum_{i=0}^{M-1} \frac{(u_{i+1} - u_i)^2}{2Gh} + 4\pi h \sum_{i=1}^{M-1} r_i \rho_i u_i. \tag{7}$$

Proposition 2.1 (Unique finite radial solution)**.** *With u_0, u_M fixed, S_h has a unique minimizer. Its interior values satisfy*

$$\frac{u_{i+1} - 2u_i + u_{i-1}}{h^2} = 4\pi G r_i \rho_i.$$

Proof. Differentiation with respect to u_i gives $(2u_i - u_{i-1} - u_{i+1})/(Gh) + 4\pi h r_i \rho_i$. Setting it to zero gives the equation. For a variation v with $v_0 = v_M = 0$, the Hessian quadratic form is

$$\frac{1}{Gh} \sum_{i=0}^{M-1} (v_{i+1} - v_i)^2.$$

Its vanishing forces v to be constant and hence zero. The finite Hessian is positive definite, so the quadratic action has one stationary point and that point is its unique minimizer. □

The boundary value $u_0 = 0$ allows a regular finite potential at the origin. The quotient u/r is used only for $r > 0$, with its continuous limit determined separately at $r = 0$.

For the normalized Gaussian density

$$\rho(r) = \frac{M_0}{(2\pi)^{3/2}\sigma^3} e^{-r^2/(2\sigma^2)}, \qquad \sigma > 0,$$

direct differentiation verifies

$$M_{\text{enc}}(r) = M_0 \left[\operatorname{erf}\left(\frac{r}{\sqrt{2}\sigma} \right) - \sqrt{\frac{2}{\pi}} \frac{r}{\sigma} e^{-r^2/(2\sigma^2)} \right],$$
$$\Phi(r) = -\frac{GM_0}{r} \operatorname{erf}\left(\frac{r}{\sqrt{2}\sigma} \right).$$

The potential tends to zero at infinity and has the finite central value $-\sqrt{2/\pi}\, GM_0/\sigma$. This removes the point-source singularity while preserving the attractive sign.

2.2 Clocks and straight-ray optical integrals

In the stipulated first-order isotropic weak-field metric,

$$ds^2 = -\left(1 + \frac{2\Phi}{c_0^2}\right) c_0^2 dt^2 + \left(1 - \frac{2\Phi}{c_0^2}\right) d\mathbf{x}^2, \qquad \left|\frac{\Phi}{c_0^2}\right| \ll 1,$$

a static clock has lapse $\mathcal{N} = \sqrt{1 + 2\Phi/c_0^2}$. The static-frequency ratio in this approximation is

$$\frac{\nu_{\text{obs}}}{\nu_{\text{emit}}} = \frac{\mathcal{N}_{\text{emit}}}{\mathcal{N}_{\text{obs}}} = 1 + \frac{\Phi_{\text{emit}} - \Phi_{\text{obs}}}{c_0^2} + O\left(\frac{\Phi^2}{c_0^4}\right).$$

For $\Phi(x, b) = -GM_0/\sqrt{x^2 + b^2}$, $b > 0$, evaluating the first-order straight-ray integrals over $[-L, L]$ gives

$$\Delta t_{\text{Shapiro}} = -\frac{2}{c_0^3}\int_{-L}^{L} \Phi(x, b)\, dx = \frac{4GM_0}{c_0^3}\,\text{asinh}\left(\frac{L}{b}\right),$$
$$\alpha(L, b) = \frac{2}{c_0^2}\left|\int_{-L}^{L} \partial_b \Phi(x, b)\, dx\right| = \frac{4GM_0 L}{c_0^2 b\sqrt{L^2 + b^2}}.$$

Here α is the magnitude of the deflection toward the source. The limit $L \to \infty$ gives $4GM_0/(c_0^2 b)$. The displayed equalities evaluate the weak-field integrals exactly. The approximation to the physical ray already entered through the metric expansion and the undeviated integration path.

3 Gowdy evolution and reconstruction of the remaining constraint

The next step is an inhomogeneous nonlinear model with two gravitational field channels. We take the archived Gowdy reduction as specified by its action and constraint densities. The results below prove the internal variational and reconstruction identities of that reduced system.

Let P, Q be real smooth fields on $I \times \mathbb{T}$, with $\theta \in \mathbb{R}/(2\pi\mathbb{Z})$ and $\gamma(\tau) = e^{-2\tau}$. Define

$$S_G = \frac{1}{2}\int_I \int_0^{2\pi} \left[P_\tau^2 - \gamma P_\theta^2 + e^{2P}(Q_\tau^2 - \gamma Q_\theta^2)\right] d\theta\, d\tau. \tag{8}$$

Time variations vanish at the endpoints of I and spatial variations are periodic. Integration by parts gives the two Euler–Lagrange expressions

$$\mathcal{E}_P = P_{\tau\tau} - \gamma P_{\theta\theta} - e^{2P}(Q_\tau^2 - \gamma Q_\theta^2),$$
$$\mathcal{E}_Q = Q_{\tau\tau} - \gamma Q_{\theta\theta} + 2(P_\tau Q_\tau - \gamma P_\theta Q_\theta). \tag{9}$$

In particular, the Q equation can first be written as $\partial_\tau(e^{2P}Q_\tau) - \partial_\theta(e^{2P}\gamma Q_\theta) = 0$. The stationary equations are $\mathcal{E}_P = \mathcal{E}_Q = 0$.

Define the densities

$$H_G = P_\tau^2 + \gamma P_\theta^2 + e^{2P}(Q_\tau^2 + \gamma Q_\theta^2),$$
$$M_G = 2(P_\tau P_\theta + e^{2P} Q_\tau Q_\theta). \tag{10}$$

The remaining reconstruction equations are $\lambda_\tau = H_G$ and $\lambda_\theta = M_G$.

Proposition 3.1 (Local integrability and global periodicity)**.** *For twice continuously differentiable P, Q,*

$$\partial_\tau M_G - \partial_\theta H_G = 2P_\theta \mathcal{E}_P + 2e^{2P} Q_\theta \mathcal{E}_Q. \tag{11}$$

For a solution on $I \times \mathbb{T}$, a single-valued periodic λ with these two derivatives exists if and only if $\int_0^{2\pi} M_G \, d\theta = 0$ at one, and hence every, time. It is unique up to an additive constant.

Proof. Differentiate the two densities. The terms with $P_\tau P_{\tau\theta}$ and $Q_\tau Q_{\tau\theta}$ cancel, leaving

$$\begin{aligned}\partial_\tau M_G - \partial_\theta H_G = {} & 2P_\theta(P_{\tau\tau} - \gamma P_{\theta\theta}) \\ & + 2e^{2P}[Q_\theta(Q_{\tau\tau} - \gamma Q_{\theta\theta}) + 2P_\tau Q_\tau Q_\theta - P_\theta Q_\tau^2 - \gamma P_\theta Q_\theta^2].\end{aligned}$$

Substitute (9) and combine the two occurrences of $\gamma P_\theta Q_\theta^2$. This gives (11). On a solution the identity vanishes, so periodicity gives $d \int M_G \, d\theta/d\tau = 0$.

If a periodic λ exists, integrating $\lambda_\theta = M_G$ forces the zero-period condition. Conversely, fix τ_0 and define on the lifted strip

$$\lambda(\tau, \theta) = \lambda_0 + \int_{\tau_0}^{\tau} H_G(s, 0)\, ds + \int_0^{\theta} M_G(\tau, \eta)\, d\eta.$$

Its θ derivative is M_G. Its τ derivative is H_G by the integrability identity and the fundamental theorem of calculus. The zero-period condition makes it periodic. The difference of two such functions has both derivatives zero, so is constant. □

This separates a local compatibility condition from a global topology condition. Satisfying the evolution equations makes the reconstruction one-form closed. A vanishing spatial period makes it the derivative of a periodic potential.

3.1 An exact polarized mode

Set $Q = 0$ and take an integer $n \geq 1$. Define the entire function

$$J_0(x) = \sum_{m=0}^{\infty} \frac{(-1)^m (x^2/4)^m}{(m!)^2}.$$

The series and its differentiated series converge uniformly on compact sets. Coefficient comparison gives $x^2 J_0'' + x J_0' + x^2 J_0 = 0$. Thus

$$P(\tau, \theta) = A J_0(n e^{-\tau}) \cos(n\theta) \tag{12}$$

is an exact polarized solution. Indeed, for $x = ne^{-\tau}$, $\partial_\tau^2 J_0(x) = x^2 J_0''(x) + x J_0'(x) = -x^2 J_0(x)$, whereas $P_{\theta\theta} = -n^2 P$. Also $M_G = 2P_\tau P_\theta$ has zero spatial integral, so periodic λ can be reconstructed.

The Taylor channels can be separated without changing the physical mode. Put $y = x^2/4$, $\mathcal{B}(y) = J_0(2\sqrt{y})$, and for an integer $r \geq 2$ define

$$\mathcal{B}_{r,j}(y) = \sum_{q=0}^{\infty} \frac{(-1)^{rq+j} y^{rq+j}}{((rq+j)!)^2}, \qquad j = 0, \ldots, r-1.$$

Then $\mathcal{B} = \sum_j \mathcal{B}_{r,j}$. With $\vartheta = y\,d/dy$, the coefficient identity

$$\vartheta^2 \mathcal{B}_{r,j} + y\mathcal{B}_{r,j-1} = 0, \qquad j \pmod r,$$

follows from $m^2/(m!)^2 = 1/((m-1)!)^2$ for $m \geq 1$. Consequently the separated Taylor channels are coupled. Their sum satisfies the Bessel equation. This decomposition resolves series order and does not introduce additional independent gravitational degrees of freedom.

4 An exact finite Gowdy constraint projection

Let $h = 2\pi/N$ and use the periodic centered derivative

$$(D_h f)_j = \frac{f_{j+1} - f_{j-1}}{2h}, \qquad (\Delta_h f)_j = \frac{f_{j+1} - 2f_j + f_{j-1}}{h^2}.$$

Here $D_h^{\mathsf{T}} = -D_h$. With $U = \dot{P}$, $V = \dot{Q}$, the raw semidiscrete vector field $f_h(\tau, z)$ for $z = (P, Q, U, V)$ is

$$\begin{aligned} \dot{P} &= U, & \dot{Q} &= V, \\ \dot{U} &= \gamma\Delta_h P + e^{2P}[V^2 - \gamma(D_h Q)^2], \\ \dot{V} &= \gamma\Delta_h Q - 2[UV - \gamma(D_h P)(D_h Q)]. \end{aligned} \tag{13}$$

Products and exponentials in these formulas are componentwise. The finite global momentum is

$$C_h(z) = 2h\sum_j [U_j (D_h P)_j + e^{2P_j} V_j (D_h Q)_j]. \tag{14}$$

It is linear in the velocity variables when the configuration P, Q is fixed.

For the ordinary Euclidean gradients, skew-adjointness of D_h gives

$$\begin{aligned} \nabla_P C_h &= -2hD_h U + 4he^{2P} V(D_h Q), \\ \nabla_Q C_h &= -2hD_h(e^{2P} V), \\ g_U := \nabla_U C_h &= 2hD_h P, & g_V := \nabla_V C_h &= 2he^{2P} D_h Q. \end{aligned}$$

Put $G_h = (0, 0, g_U, g_V)$ and $d_h = \|g_U\|_2^2 + \|g_V\|_2^2$. On the open region $d_h > 0$, define

$$\Lambda_h = \frac{\nabla C_h \cdot f_h}{d_h}, \qquad \hat{f}_h = f_h - \Lambda_h G_h.$$

The chain rule proves the exact tangency identity

$$\nabla C_h \cdot \hat{f}_h = \nabla C_h \cdot f_h - \Lambda_h d_h = 0.$$

Every differentiable solution that remains in this region therefore keeps C_h constant.

Proposition 4.1 (Exact endpoint momentum)**.** *Start from z_n and let a provisional step produce $\tilde{z} = (\widetilde{P}, \widetilde{Q}, \widetilde{U}, \widetilde{V})$. Form $\tilde{g}_U, \tilde{g}_V$ at its configuration and assume $\tilde{d}_h > 0$. The correction*

$$\mu = \frac{C_h(\tilde{z}) - C_h(z_n)}{\tilde{d}_h}, \qquad z_{n+1} = (\widetilde{P}, \widetilde{Q}, \widetilde{U} - \mu\tilde{g}_U, \widetilde{V} - \mu\tilde{g}_V)$$

satisfies $C_h(z_{n+1}) = C_h(z_n)$ exactly. It is the correction of smallest Euclidean velocity norm at the fixed configuration.

Proof. Linearity in the velocities gives $C_h(z_{n+1}) = C_h(\tilde{z}) - \mu\tilde{d}_h$. For minimality, write $g = (\tilde{g}_U, \tilde{g}_V)$ and $b = C_h(z_n) - C_h(\tilde{z})$. Every admissible velocity displacement v satisfies $g \cdot v = b$. Cauchy–Schwarz implies $\|v\| \geq |b|/\|g\|$, with equality only at $v = bg/\|g\|^2$. This is the displayed correction. □

Proposition 4.2 (Fixed-mesh second-order time accuracy)**.** *Fix h. Suppose a smooth exact solution of $\dot{z} = \hat{f}_h(\tau, z)$ and its numerical steps remain in a compact regular neighborhood where $d_h \geq d_0 > 0$. Assume the mixed time and state derivatives of $\hat{f}_h$ through order three and the state derivatives of C_h through order four are bounded there. Apply the endpoint correction of Proposition 4.1 after explicit midpoint. Then its local error is $O(\Delta\tau^3)$ and its global error on a fixed time interval is $O(\Delta\tau^2)$.*

Proof. The explicit midpoint expansion agrees with the Taylor expansion of the exact flow through order two. Its provisional endpoint differs from the exact endpoint by $O(\Delta\tau^3)$. Since the exact flow keeps C_h constant and C_h has bounded first derivatives, the provisional constraint defect is also $O(\Delta\tau^3)$. The lower bound on the denominator and the upper bound on g make the correction $O(\Delta\tau^3)$. Thus the corrected local error has the same order. Under the stated derivative bounds the corrected step map has Lipschitz constant at most $1 + L\Delta\tau$ for small steps. If e_n is its error against the exact solution, then

$$e_{n+1} \leq (1 + L\Delta\tau)e_n + M\Delta\tau^3.$$

Iterating this inequality for $n\Delta\tau \leq T$ bounds the error by $MTe^{LT}\Delta\tau^2$ when $e_0 = 0$. □

These are regular-chart results. For a spatially homogeneous configuration, $g_U = g_V = 0$, and the displayed multiplier is undefined even though $C_h = 0$. Constants in the accuracy theorem can depend on h. A claim about simultaneous spatial and temporal refinement additionally needs bounds uniform in h.

The local compatibility diagnostic is a separate quantity. Define $H_{G,h}, M_{G,h}$ by replacing spatial derivatives in (10) with D_h. If $\mathcal{E}_{P,h} = \dot{U} - \dot{U}_{\text{raw}}$ and $\mathcal{E}_{Q,h} = \dot{V} - \dot{V}_{\text{raw}}$, then

$$\mathcal{J}_h = \partial_\tau M_{G,h} - D_h H_{G,h} - 2(D_h P)\mathcal{E}_{P,h} - 2e^{2P}(D_h Q)\mathcal{E}_{Q,h}$$

subtracts the explicit constraint-force contribution. Exact global momentum preservation does not force $\mathcal{J}_h$ to vanish at each site. Its refinement behavior tests whether the modified evolution approaches the local continuum reconstruction identity. The archived second-order numerical observations are experimental evidence for that behavior, with their tested range recorded in the companion.

5 Finite linearized ADM gravity

The symmetry-reduced model can now be compared with a finite three-dimensional linearized constraint system. Its exact algebraic closure rests on commuting linear derivatives. Nonlinear products will introduce an additional requirement in the next section.

5.1 Finite calculus and the quadratic action

Let V_h be a finite real inner-product space. Let $\delta_1, \delta_2, \delta_3$ be commuting skew-adjoint operators on it:

$$\delta_i\delta_j = \delta_j\delta_i, \qquad \langle \delta_i X, Y\rangle_h = -\langle X, \delta_i Y\rangle_h, \qquad \delta^2 = \sum_i \delta_i^2.$$

The periodic centered differences in three coordinate directions provide one example. All tensor inner products below sum over both indices, using the Frobenius inner product on symmetric tensors.

Let $h_{ij} = h_{ji}$ and $\pi_{ij} = \pi_{ji}$ take values in V_h. Repeated spatial indices are summed. Put $h_{\mathrm{tr}} = h_{kk}$, $\pi_{\mathrm{tr}} = \pi_{kk}$ and define

$$\begin{aligned} C &= \delta_i\delta_j h_{ij} - \delta^2 h_{\mathrm{tr}}, & C_i &= -2\delta_j \pi_{ij}, \\ \mathcal{G}_{ij}(h) &= \frac{1}{2}[\delta_k\delta_i h_{jk} + \delta_k\delta_j h_{ik} - \delta^2 h_{ij} - \delta_i\delta_j h_{\mathrm{tr}} - \delta_{ij} C]. \end{aligned} \tag{15}$$

Lemma 5.1 (Linearized finite Bianchi and adjoint identities)**.** *The spatial operator obeys $\delta_j \mathcal{G}_{ij}(h) = 0$ and is self-adjoint on symmetric tensors.*

Proof. Twice its divergence expands as

$$\begin{aligned} 2\delta_j\mathcal{G}_{ij} = {} & \delta_i\delta_j\delta_k h_{jk} + \delta^2\delta_k h_{ik} - \delta^2\delta_j h_{ij} - \delta_i\delta^2 h_{\mathrm{tr}} \\ & - \delta_i\delta_k\delta_l h_{kl} + \delta_i\delta^2 h_{\mathrm{tr}}. \end{aligned}$$

The terms cancel in three pairs. For self-adjointness, write $B_i(a) = \delta_j a_{ij}$ and $A(a) = \delta_i\delta_j a_{ij}$. Two applications of the adjoint identity give

$$\begin{aligned} 2\langle a, \mathcal{G}(b)\rangle_h = {} & -2\langle B(a), B(b)\rangle_h + \sum_k \langle \delta_k a, \delta_k b\rangle_h \\ & - \langle A(a), b_{\mathrm{tr}}\rangle_h - \langle a_{\mathrm{tr}}, A(b)\rangle_h - \sum_k \langle \delta_k a_{\mathrm{tr}}, \delta_k b_{\mathrm{tr}}\rangle_h . \end{aligned}$$

This expression is symmetric in a, b. □

Set

$$T_h = \langle \pi, \pi\rangle_h - \tfrac{1}{2}\langle \pi_{\mathrm{tr}}, \pi_{\mathrm{tr}}\rangle_h, \qquad U_h = \tfrac{1}{2}\langle h, \mathcal{G}(h)\rangle_h .$$

With lapse perturbation n and shift perturbation n_i, the action is

$$S_h^{(2)} = \int \Big[\big\langle \pi, \dot h\big\rangle_h - T_h - U_h + \langle n, C\rangle_h - \langle n_i, C_i\rangle_h\Big]\, dt. \tag{16}$$

Variation of n, n_i imposes $C = C_i = 0$. The remaining equations are

$$\begin{aligned} \dot h_{ij} &= 2\pi_{ij} - \delta_{ij}\pi_{\mathrm{tr}} + \delta_i n_j + \delta_j n_i, \\ \dot\pi_{ij} &= -\mathcal{G}_{ij}(h) + \delta_i\delta_j n - \delta_{ij}\delta^2 n. \end{aligned} \tag{17}$$

The kinetic gradient is $2\pi_{ij} - \delta_{ij}\pi_{\mathrm{tr}}$. Lemma 5.1 makes the potential gradient $\mathcal{G}(h)$. The lapse gradient follows by two spatial integrations by parts. Integrating the shift term once gives $-\langle \delta_i n_j + \delta_j n_i, \pi_{ij}\rangle_h$ in the Lagrangian, fixing the sign in the first equation.

5.2 Constraint propagation and the canonical brackets

Theorem 5.2 (Exact finite linearized closure)**.** *For every solution of* (17)*,*

$$\dot{C} = -\delta_i C_i, \qquad \dot{C}_i = 0.$$

Thus initially vanishing constraints remain zero. With the canonical Frobenius bracket and smearings $C[n] = \langle n, C\rangle_h$*,* $D[\xi] = \langle \xi_i, C_i\rangle_h$*, all three brackets* $\{C[n], C[m]\}$*,* $\{D[\xi], D[\eta]\}$*, and* $\{C[n], D[\xi]\}$ *vanish identically.*

Proof. The trace of the momentum part of $\dot{h}$ is $-\pi_{\mathrm{tr}}$. Its contribution to $\dot{C}$ is therefore $2\delta_i\delta_j\pi_{ij} = -\delta_i C_i$. The shift contribution is

$$\delta_i\delta_j(\delta_i n_j + \delta_j n_i) - \delta^2(2\delta_k n_k) = 0.$$

For $\dot{C}_i$, the divergence of $\mathcal{G}$ vanishes by Lemma 5.1, and $\delta_j(\delta_i\delta_j n - \delta_{ij}\delta^2 n) = 0$. This proves the propagation equations and their stated implication.

For the bracket, use $\{F, G\} = \langle \nabla_h F, \nabla_\pi G\rangle_h - \langle \nabla_\pi F, \nabla_h G\rangle_h$. Two scalar constraints depend only on h and two momentum constraints depend only on π, so their brackets vanish. The gradients for the mixed bracket are

$$(\nabla_h C[n])_{ij} = \delta_i\delta_j n - \delta_{ij}\delta^2 n, \qquad (\nabla_\pi D[\xi])_{ij} = \delta_i\xi_j + \delta_j\xi_i.$$

Their contraction vanishes after one integration by parts because the first tensor is divergence-free. □

The momentum generator changes the metric perturbation by $h_{ij} \mapsto h_{ij} + \delta_i\xi_j + \delta_j\xi_i$. The corresponding change of the scalar constraint is $2\delta^2\delta_k\xi_k - 2\delta^2\delta_k\xi_k = 0$. No product rule is used in these cancellations.

For a partition of fine spatial sites into coarse regions A, the nonnegative diagnostic

$$E_A^{\mathrm{con}} = \sum_{x\in A} h^3 \left(|C_x|^2 + \sum_i |C_{i,x}|^2 \right)$$

adds exactly to the full fine diagnostic. This is an exact partition identity. It does not assert that differentiating a coarse field reconstructs its fine constraint values.

5.3 Tensor waves, positive reduced energy, and resolved dispersion

Choose $n = n_i = 0$ and impose the transverse and traceless conditions

$$\delta_j h_{ij} = 0, \quad h_{\mathrm{tr}} = 0, \qquad \delta_j\pi_{ij} = 0, \quad \pi_{\mathrm{tr}} = 0.$$

They are preserved by the evolution because

$$\mathcal{G}_{ij}(h) = -\tfrac{1}{2}\delta^2 h_{ij}, \qquad \dot{h} = 2\pi, \qquad \dot{\pi} = \tfrac{1}{2}\delta^2 h, \qquad \ddot{h} = \delta^2 h.$$

On this reduced sector the energy is

$$H_{\mathrm{TT}} = \|\pi\|_h^2 + \frac{1}{4}\sum_k \|\delta_k h\|_h^2 \geq 0.$$

Its derivative is zero: substituting the two first-order equations reduces it to $\langle \pi, \delta^2 h\rangle_h + \sum_k \langle \delta_k h, \delta_k \pi\rangle_h = 0$.

For a Fourier mode on the periodic centered lattice, let $\tilde{k}_i = \sin(k_i h)/h$. If $\tilde{\mathbf{k}} \neq 0$, define

$$\Pi_{ij} = \delta_{ij} - \frac{\tilde{k}_i \tilde{k}_j}{|\tilde{\mathbf{k}}|^2}, \qquad \Lambda^{\mathrm{TT}}_{ij,kl} = \tfrac{1}{2}(\Pi_{ik}\Pi_{jl} + \Pi_{il}\Pi_{jk} - \Pi_{ij}\Pi_{kl}).$$

For a symmetric tensor X, its image is $\Pi X \Pi - \Pi\,\mathrm{tr}(\Pi X)/2$. It is transverse and traceless because $\Pi^2 = \Pi$ and $\mathrm{tr}\,\Pi = 2$. Applying the map again leaves it unchanged. On the transverse two-plane, symmetric tensors have dimension three and the trace condition removes one dimension. Thus the projector has rank two.

The resulting frequency obeys

$$\omega_h^2 = \sum_i \frac{\sin^2(k_i h)}{h^2}. \tag{18}$$

For fixed $\mathbf{k} \neq 0$ as $h \to 0$,

$$\omega_h^2 = |\mathbf{k}|^2 - \frac{h^2}{3}\sum_i k_i^4 + O(h^4), \qquad \omega_h = |\mathbf{k}| - \frac{h^2}{6|\mathbf{k}|}\sum_i k_i^4 + O(h^4).$$

These follow by expanding the sine and then the positive square root. The regime keeps $\mathbf{k}$ fixed. It is not uniform up to the Nyquist boundary.

The spatial operator in (18) is the sum of the squared centered first derivatives. It differs from the nearest-neighbor second-difference operator used for the raw Gowdy equations. For even site counts, an alternating sequence has centered derivative zero. Such checkerboard modes, as well as constant modes, belong to the derivative kernel. At a zero symbol the above TT projector is undefined and the rank-two argument does not apply. Exact constraint propagation consequently coexists with a finite spatial kernel that a physical continuum comparison must address.

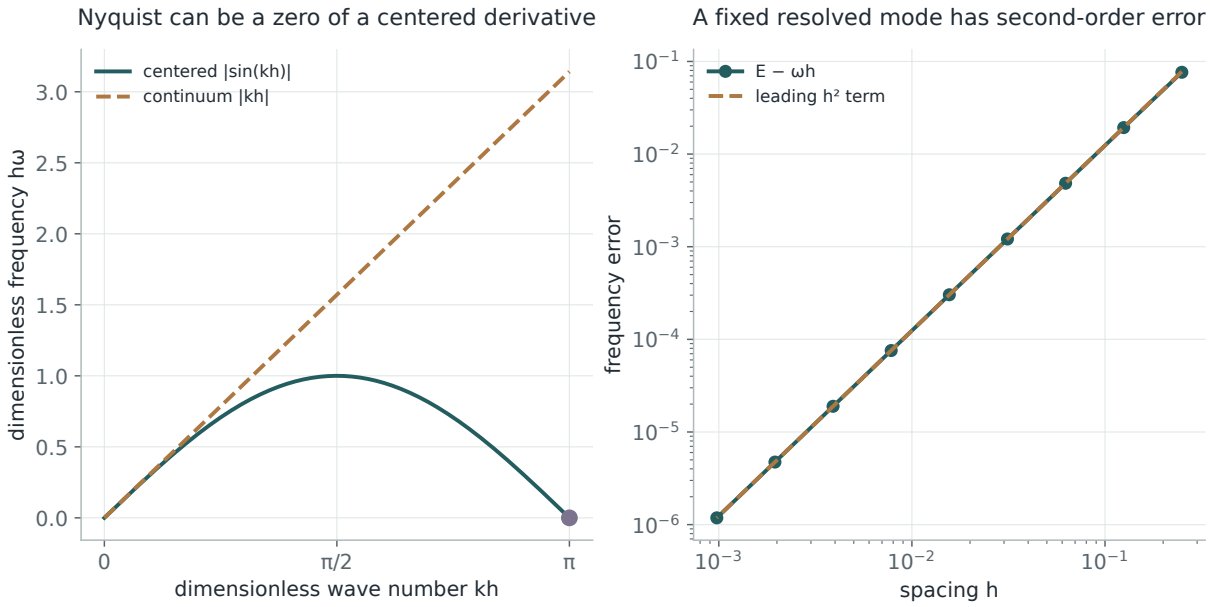


Figure 2. The centered derivative has a Nyquist zero, while a fixed resolved mode approaches its continuum frequency with second-order error. The left panel is the one-dimensional symbol; the right panel tests a fixed mode under refinement.

6 The nonlinear pointwise obstruction and its algebraic alternative

The linearized cancellations use commuting derivatives but never ask them to differentiate a product. A nonlinear diffeomorphism generator does. The following scalar model isolates that issue without relying on the more involved gravitational tensor constraints.

Let a periodic real lattice have canonical brackets $\{q_a, p_b\}_h = \delta_{ab}/h$ and a skew matrix D. Define

$$\mathcal{D}_h[\xi] = h\sum_a \xi_a p_a (Dq)_a, \qquad \mathfrak{L}_D(f,g) = D(fg) - (Df)g - f(Dg).$$

All products on this lattice are pointwise.

Proposition 6.1 (Exact nonlinear defect identity). *For arbitrary finite arrays q, p, ξ, η,*

$$\begin{aligned}\{\mathcal{D}_h[\xi], \mathcal{D}_h[\eta]\}_h &= \mathcal{D}_h[\xi D\eta - \eta D\xi] \\ &\quad + h\sum_a p_a[\xi_a \mathfrak{L}_D(\eta, Dq)_a - \eta_a \mathfrak{L}_D(\xi, Dq)_a].\end{aligned} \tag{19}$$

Proof. Coordinate differentiation gives $\partial_p \mathcal{D}_h[\xi] = h\xi Dq$ and $\partial_q \mathcal{D}_h[\xi] = -hD(\xi p)$. Insert these in the canonical bracket and use $D^{\mathsf{T}} = -D$:

$$\{\mathcal{D}_h[\xi], \mathcal{D}_h[\eta]\}_h = h\sum_a p_a\left[\xi_a D(\eta Dq)_a - \eta_a D(\xi Dq)_a\right].$$

Expanding each derivative using the definition of $\mathfrak{L}_D$ cancels the two $\xi\eta D^2 q$ terms. The remaining product-rule terms give the first term of (19), and the two defect terms give the second. □

Theorem 6.2 (Finite pointwise derivation obstruction)**.** *A linear derivation on the complete pointwise algebra $\mathbb{C}^N$ is zero. Thus a nonzero finite pointwise derivative cannot obey the undeformed Leibniz rule on every pair of arrays.*

Proof. Let e_a be the coordinate idempotents. If D is a derivation, then $D(e_a) = D(e_a^2) = 2e_aD(e_a)$. At coordinate a this forces the component to be twice itself, hence zero. At every other coordinate the right side is zero. Therefore $D(e_a) = 0$ for every a. Linearity and the coordinate basis imply $D = 0$. □

The theorem rules out a universal ordinary derivation on the full finite pointwise algebra. For a particular generator, one must still inspect the contracted defect in (19). A nonzero Leibniz defect does not by itself imply that every contraction is nonzero.

Example 6.3 (An exact bracket failure)**.** Use $N = 5$, $h = 1$, and $(Df)_a = (f_{a+1} - f_{a-1})/2$. Take $q = e_0$, $p = \xi = e_1$, and $\eta = e_2$. Then $(Dq)_1 = -1/2$, $(D\eta)_1 = 1/2$, and $\eta Dq = 0$. The direct bracket in the preceding proof is zero, while

$$\mathcal{D}_h[\xi D\eta - \eta D\xi] = -\tfrac{1}{4}.$$

The defect is therefore exactly $+1/4$. This finite witness shows that the naive centered scalar constraint does fail to close universally in the proposed undeformed form.

6.1 Smooth resolved arrays and the matrix corner route

For samples of sufficiently smooth periodic functions, the centered derivative obeys

$$D_hf = f' + \frac{h^2}{6}f''' + O(h^4).$$

If f, g have uniformly bounded derivatives through order five, its Leibniz defect consequently satisfies

$$\mathfrak{L}_{D_h}(f, g) = \frac{h^2}{2}(f''g' + f'g'') + O(h^4). \tag{20}$$

For fixed smooth p, ξ, η, q, with sufficiently many additional bounded derivatives of q to control D_hq, the contraction in (19) is $O(h^2)$ on a fixed periodic interval. Each summand is uniformly $O(h^2)$ and hN is the interval length. The coordinate-idempotent witness above is outside this uniformly smooth regime.

A spectral derivative has an exact product rule on trigonometric polynomials when every product remains in the represented band without aliasing. Repeated products require further headroom. These restricted domains differ from the complete pointwise algebra. The later resolved constraint theorem includes explicit domain bounds.

For the full matrix algebra, choose a Hermitian matrix K and set $\delta_KX = i[K, X]$. Direct multiplication gives

$$\delta_K(XY) = (\delta_KX)Y + X(\delta_KY). \tag{21}$$

It also obeys $\delta_K(X^*) = (\delta_KX)^*$, $\mathrm{Tr}\,\delta_KX = 0$, and is skew-adjoint for the Hilbert–Schmidt inner product. The last identity follows by moving K through the trace. These properties supply an exact finite product calculus on a noncommutative algebra.

Nonlinear algebra and regular reduction

7 Nonlinear closure and the return of discarded detail

The nonlinear construction has two distinct tasks. A fine theory must first have a consistent gauge differential. Its passage to a retained resolution must then preserve that differential. The second task can be solved conditionally without solving the first. This distinction explains the sequence of the chapters: an obstruction to direct lattice closure is followed by a transfer mechanism, a resolved vacuum example, a real Lorentzian parent, and a local canonical reduction. Each step changes the mathematical setting. An earlier obstruction is not a refutation of a later construction on a different domain, and the later construction does not make the earlier hypotheses disappear.

7.1 An exact matrix momentum sector

Let the fine canonical coordinates be Hermitian matrix pairs (Z^A, P_A), with real components and the trace bracket

$$\{F, G\} = \sum_A \mathrm{Tr}\left(\nabla_{Z^A} F \nabla_{P_A} G - \nabla_{P_A} F \nabla_{Z^A} G\right). \tag{22}$$

The trace gradients are defined by real variations. For a Hermitian parameter Λ, put

$$D[\Lambda] = i \sum_A \mathrm{Tr}(P_A[\Lambda, Z^A]). \tag{23}$$

This real functional generates simultaneous conjugation:

$$\{Z^A, D[\Lambda]\} = i[\Lambda, Z^A], \qquad \{P_A, D[\Lambda]\} = i[\Lambda, P_A]. \tag{24}$$

Indeed, cyclic variation gives $\nabla_P D[\Lambda] = i[\Lambda, Z]$ and $\nabla_Z D[\Lambda] = i[P, \Lambda]$ for one pair.

Proposition 7.1 (Finite inner momentum closure)**.** *With these conventions,*

$$\{D[\Lambda], D[\Gamma]\} = D[i[\Lambda, \Gamma]]. \tag{25}$$

If a matrix density $\mathcal{E}$ obeys $\{\mathcal{E}, D[\Lambda]\} = i[\Lambda, \mathcal{E}]$, then $H[N] = \mathrm{Tr}(N\mathcal{E})$ satisfies

$$\{D[\Lambda], H[N]\} = H[i[\Lambda, N]]. \tag{26}$$

Proof. For one pair, substitution of the two cyclic gradients yields

$$\begin{aligned} \{D[\Lambda], D[\Gamma]\} &= \mathrm{Tr}\left(-[P, \Lambda][\Gamma, Z] + [\Lambda, Z][P, \Gamma]\right) && (27)\\ &= \mathrm{Tr}\left(P(-[\Lambda, [\Gamma, Z]] + [\Gamma, [\Lambda, Z]])\right) && (28)\\ &= -\mathrm{Tr}(P[[\Lambda, \Gamma], Z]) = D[i[\Lambda, \Gamma]]. && (29) \end{aligned}$$

Summation handles all pairs. For the second statement, antisymmetry and cyclicity give

$$\{D[\Lambda], H[N]\} = -i\,\mathrm{Tr}(N[\Lambda, \mathcal{E}]) = \mathrm{Tr}(i[\Lambda, N]\mathcal{E}). \tag{30}$$

□

The closed finite object is the full inner algebra. To identify a continuum spatial diffeomorphism, one must additionally construct parameters $\Lambda_h[\xi]$ whose action approximates the Lie derivative on resolved observables. An arbitrary subset of matrices chosen to resemble first derivatives need not be closed under commutators. Equation (25) does not establish that extra geometric identification.

7.2 What direct projection loses

Let $(\mathcal{C}_F, s_F)$ be a graded fine complex with $s_F^2 = 0$. Let $P : \mathcal{C}_F \to \mathcal{C}_C$ and $I : \mathcal{C}_C \to \mathcal{C}_F$ be degree-zero maps satisfying $PI = \mathrm{id}_C$, and set

$$Q = \mathrm{id}_F - IP, \qquad \bar{s}_C = P s_F I. \tag{31}$$

Here I is the inclusion map, whereas id denotes an identity operator. Insertion of $IP = \mathrm{id}_F - Q$ proves the exact identity

$$\bar{s}_C^2 = P s_F (IP) s_F I = -P s_F Q s_F I. \tag{32}$$

The lost path starts in retained data, enters a discarded channel, and returns to retained data. Its sign and order are fixed by the formula. There is no reason for it to vanish merely because the fine differential is nilpotent.

For example, the finite transfer test has the naive coarse block

$$\bar{s}_C = \begin{pmatrix} -1/25 & -1/25 & -26/625 \\ 1 & 0 & 1/25 \\ 0 & 1/25 & 0 \end{pmatrix}. \tag{33}$$

Its square is nonzero, as its $(2,2)$ entry is $-1/25$. The transferred block in that same test is

$$s_C = \begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & 1/25 \\ 0 & 0 & 0 \end{pmatrix}, \qquad s_C^2 = 0. \tag{34}$$

This small rational example illustrates the distinction between a projection and a transfer. The matrices by themselves are an algebraic diagnostic. They do not define gravitational constraints, a symplectic BFV theory, or a continuum model.

7.3 A precise conditional transfer theorem

Split $s_F = s_0 + \delta$, where $s_0^2 = 0$ and δ has degree one. Suppose $(\mathcal{C}_C, s_{0,C})$ is a coarse complex and the maps are chain maps:

$$s_0 I = I s_{0,C}, \qquad P s_0 = s_{0,C} P. \tag{35}$$

Assume a degree-minus-one homotopy H satisfies

$$\begin{aligned} PI = \mathrm{id}_C, \qquad & \mathrm{id}_F - IP = s_0 H + H s_0, \\ PH = 0, \qquad HI = 0, \qquad & H^2 = 0. \end{aligned} \tag{36}$$

Finally, require $(s_0+\delta)^2=0$ and an inverse of $\mathrm{id}_F+H\delta$. The inverse may exist in a declared operator topology, or in a complete filtration in which the Neumann series makes formal sense. A formal inverse is a formal conclusion, not an analytic convergence theorem.

Theorem 7.2 (Transfer of a differential)**.** *Under these assumptions, define*

$$I_\delta=(\mathrm{id}_F+H\delta)^{-1}I, \qquad s_C=s_{0,C}+P\delta I_\delta. \tag{37}$$

Then $PI_\delta=\mathrm{id}_C$, $(s_0+\delta)I_\delta=I_\delta s_C$, *and* $s_C^2=0$.

Proof. The inverse equation is $I_\delta=I-H\delta I_\delta$. Application of P proves $PI_\delta=PI=\mathrm{id}_C$. Let

$$K=(s_0+\delta)I_\delta-I_\delta s_C. \tag{38}$$

Using the chain-map identity and $\mathrm{id}_F-IP=s_0H+Hs_0$ in this expression gives

$$K=H\left(s_0\delta I_\delta+\delta I_\delta s_{0,C}+\delta I_\delta P\delta I_\delta\right). \tag{39}$$

Nilpotency of the fine differential gives $s_0\delta+\delta s_0+\delta^2=0$. The parenthesis therefore equals $-\delta K$. Hence $(\mathrm{id}_F+H\delta)K=0$, and invertibility gives $K=0$. Applying the fine differential to the intertwining identity yields $0=I_\delta s_C^2$. The left inverse P then gives $s_C^2=0$. □

Expansion of the inverse exhibits the returned detail:

$$\begin{aligned} s_C=s_{0,C}&+P\delta I-P\delta H\delta I\\ &+P\delta H\delta H\delta I-\cdots. \end{aligned} \tag{40}$$

The first correction describes one propagation through H between two nonlinear vertices. More interactions require further terms. Truncating this series generally changes the nilpotency statement.

There is a separate step from this theorem to a nonlinear BFV charge. On a canonical graded phase space, introduce odd ghosts c^A of ghost number one and their odd conjugate momenta ϖ_A of ghost number minus one. A charge of ghost number one has the form

$$\Omega_F=c^AC_A-\tfrac{1}{2}c^Ac^BU_{AB}{}^C\varpi_C+\Omega_F^{(\geq 3)}. \tag{41}$$

The superscript counts ghost factors, not total ghost number. Higher terms can have further ghost momenta and still have ghost number one. They carry the additional identities needed for field-dependent, open, or reducible algebras. They cannot be suppressed just because the displayed constraint brackets close.

If $\{\Omega_F,\Omega_F\}=0$, then $s_F=\{\Omega_F,\cdot\}$ is a differential. To transfer a nonlinear Hamiltonian gauge structure, the contraction must also lift compatibly to the completed nonlinear coalgebra and respect its cyclic or symplectic structure. On a local exact canonical chart, a resulting Hamiltonian vector field can have a charge Ω_C. Its square then has Hamiltonian $\frac{1}{2}\{\Omega_C,\Omega_C\}$, which vanishes when the canonical center has no ghost-number-two constants. These are additional assumptions. Connectedness alone does not turn every symplectic vector field into a global Hamiltonian vector field.

7.4 A quantitative return bound

For bounded maps on declared normed spaces, assume

$$\|H\delta\| \le \rho < 1, \qquad \|Q\delta I\| \le Ch^r. \tag{42}$$

The resolvent identity and $H = HQ$ give

$$s_C - (s_{0,C} + P\delta I) = -P\delta(\mathrm{id}_F + H\delta)^{-1}H\delta I, \tag{43}$$

$$\|s_C - (s_{0,C} + P\delta I)\| \le \frac{\|P\delta\|\,\|H\|}{1-\rho}\,Ch^r. \tag{44}$$

To make this an approximation theorem along a refinement sequence, the prefactor must be bounded uniformly in h. An unbounded detail Green operator can defeat the conclusion even when $Q\delta I$ is small. The estimate controls the stated differential maps. It does not by itself prove convergence of constrained trajectories, physical observables, or a quantum measure.

8 A resolved nonlinear vacuum algebra

8.1 A filtered algebra with room for products

On $\mathbb{T}^3 = (\mathbb{R}/2\pi\mathbb{Z})^3$, let

$$\mathcal{T}_K = \left\{ \sum_{\|k\|_\infty \le K} \hat{f}_k e^{ik\cdot x} \right\}. \tag{45}$$

The algebraic relations are

$$\mathcal{T}_K\mathcal{T}_L \subset \mathcal{T}_{K+L}, \qquad \partial_a\mathcal{T}_K \subset \mathcal{T}_K. \tag{46}$$

Thus $\mathcal{T}_K$ is a resolved space, while $\bigcup_K \mathcal{T}_K$ is the product-closed differential algebra. Let S_N sample on an N^3 periodic grid, δ_a be its Fourier derivative, and $\tau_N = N^{-3}\sum_x$ its normalized trace.

If $f_j \in \mathcal{T}_{K_j}$ and $\sum_j K_j < N/2$, the product stays strictly below Nyquist in each coordinate. Sampling is injective on this support and gives

$$\delta_a \prod_j S_N f_j = S_N \partial_a \left(\prod_j f_j \right), \tag{47}$$

$$\tau_N \left(\prod_j S_N f_j \right) = (2\pi)^{-3} \int_{\mathbb{T}^3} \prod_j f_j \, \mathrm{d}^3x. \tag{48}$$

The Leibniz rule and trace integration by parts consequently hold for the declared products. They do not hold for arbitrary sampled pointwise products. This is how the resolved construction avoids the finite pointwise derivation obstruction: it retains enough detail instead of demanding a derivation on the whole finite function algebra.

8.2 Constraints and their normalizations

Take connection coefficients A_a^i, densitized triads E_i^a, internal parameters λ^i, shifts ξ^a, and densitized lapse n in $\mathcal{T}_K$. Set the canonical normalization to one and define

$$F_{ab}^i = \delta_a A_b^i - \delta_b A_a^i + \epsilon^i{}_{jk} A_a^j A_b^k, \tag{49}$$

$$\mathcal{G}_i = \delta_a E_i^a + \epsilon_{ij}{}^k A_a^j E_k^a, \tag{50}$$

$$\mathcal{V}_a = F_{ab}^i E_i^b, \qquad \mathcal{D}_a = \mathcal{V}_a - A_a^i \mathcal{G}_i, \tag{51}$$

$$\mathcal{H} = \tfrac{1}{2}\epsilon^{ij}{}_k E_i^a E_j^b F_{ab}^k + \Lambda \det E. \tag{52}$$

Here $\det E = \frac{1}{6}\epsilon_{abc}\epsilon^{ijk} E_i^a E_j^b E_k^c$. The smeared generators are

$$G[\lambda] = \tau_N(\lambda^i \mathcal{G}_i), \qquad D[\xi] = \tau_N(\xi^a \mathcal{D}_a), \tag{53}$$

$$V[\xi] = \tau_N(\xi^a \mathcal{V}_a), \qquad H[n] = \tau_N(n\mathcal{H}). \tag{54}$$

On the full grid canonical phase space the bracket is

$$\{F, G\}_N = N^3 \sum_{x,a,i} \left(\frac{\partial F}{\partial A_a^i(x)} \frac{\partial G}{\partial E_i^a(x)} - \frac{\partial F}{\partial E_i^a(x)} \frac{\partial G}{\partial A_a^i(x)} \right). \tag{55}$$

The factor N^3 matches the normalized trace in the action. The resolved identities below are evaluated at band-limited data in this phase space. Replacing this bracket by a bracket that first projects every gradient back to $\mathcal{T}_K$ is a different operation and can reintroduce the leakage of equation (32).

Put

$$(\lambda \times \mu)^i = \epsilon^i{}_{jk} \lambda^j \mu^k, \tag{56}$$

$$[\xi, \eta]^a = \xi^b \delta_b \eta^a - \eta^b \delta_b \xi^a, \tag{57}$$

$$\mathcal{L}_\xi^{(-1)} n = \xi^a \delta_a n - n \delta_a \xi^a, \tag{58}$$

$$S^a(n, m) = E_i^a E^{bi} (m \delta_b n - n \delta_b m). \tag{59}$$

The lapse convention fixes the sign and density weight of this last expression. It should not be compared with a differently ordered normal bracket without translating conventions.

Proposition 8.1 (The resolved algebraic transfer step)**.** *Assume the continuum polynomial constraint identities in this convention and* $16K < N$*. Their full-grid sampled evaluations obey*

$$\{G[\lambda], G[\mu]\}_N = G[\lambda \times \mu], \tag{60}$$

$$\{D[\xi], G[\lambda]\}_N = G[\xi^a \delta_a \lambda], \tag{61}$$

$$\{D[\xi], D[\eta]\}_N = D[[\xi, \eta]], \tag{62}$$

$$\{G[\lambda], H[n]\}_N = 0, \tag{63}$$

$$\{D[\xi], H[n]\}_N = H[\mathcal{L}_\xi^{(-1)} n], \tag{64}$$

$$\{H[n], H[m]\}_N = V[S(n, m)] = D[S(n, m)] + G[S^a(n, m) A_a]. \tag{65}$$

Proof. Including the smearing, the Gauss, improved momentum, and Hamiltonian functionals contain at most three, four, and five band-K factors. Their canonical gradients contain at most two, three, and four. A bracket therefore contains at most eight factors. The hypothesis gives $8K < N/2$. Equation (48)

identifies every product, derivative, integration by parts, and zero-mode extraction in the continuum derivation with its sampled value. The continuum identities transfer term by term. The cosmological term is gauge and diffeomorphism invariant. Its cross terms in the Hamiltonian bracket cancel by antisymmetry. The identity $V[S] = D[S] + G[S^a A_a]$ follows directly from the definition of $\mathcal{D}_a$. □

The proposition isolates the spectral transfer argument. The continuum connection calculation is its stated input. Neither a floating-point test nor a source reference can replace that input. The construction is an ordinary mathematical result rather than a certified Lean theorem.

8.3 What the closure tower does and does not preserve

At a resolved point, a Hamiltonian vector field can immediately generate higher bands. In particular, the fixed space $\mathcal{T}_K$ is not generally invariant under nonlinear evolution. The valid procedure enlarges the representation before the newly generated support aliases. Fourier zero padding satisfies

$$I_{L\to M} I_{K\to L} = I_{K\to M}. \tag{66}$$

When the larger band and its grid resolve every generated product, the embedded constraints and their bracket values agree. This is compatibility of finite algebraic evaluations. A tower of such evaluations does not yet supply a uniform interval of existence, a bound on band growth, or a continuum evolution theorem.

The source certificate tests the algebra at $N = 18$, $K = 1$, where $16K < N$. Its bracket residuals are of order 10^{-17}. A band-three negative control violates the hypothesis and produces a Hamiltonian-bracket residual near 3.3×10^{-8}. These are useful diagnostics of the implementation and of the support boundary. A finite test is not a proof for all admissible fields.

The Lorentzian limitation of this self-dual chart is also explicit. A single Lorentzian self-dual connection is complex, with a physical sector requiring real E and $A + \overline{A} = 2\Gamma(E)$. The source chapter does not construct a finite step preserving that sector. The real parent below addresses reality by changing variables and retaining both chiral sectors, rather than by declaring this missing step solved inside the single complex chart.

8.4 The separate Gowdy BFV target

The Gowdy target chapter changes from τ to areal time $t = e^{-\tau} > 0$. Its wave-map Hamiltonian is

$$\mathcal{H}_t = \frac{\Pi_P^2 + e^{-2P}\Pi_Q^2}{2t} + \frac{t}{2}(P_\theta^2 + e^{2P} Q_\theta^2). \tag{67}$$

The physical Hamiltonian density at fixed t and the tangential density $\mathcal{D} = \Pi_P P_\theta + \Pi_Q Q_\theta$ have the continuum bracket pattern of two deformation generators with densitized smearings. This algebraic target does not make $\mathcal{H}_t$ a vanishing gravitational constraint. Areal time is already a gauge choice.

A gravitational BFV construction must begin with the ungauge-fixed reduced Einstein action, retain its clock or embedding variables, derive its constraints, and recover the displayed physical Hamiltonian after gauge fixing. The source lists fine nilpotency, finite transfer, refinement consistency, continuum recovery, and quantum measure control as successive gates. The resolved vacuum example supplies a different classical input and does not close those Gowdy-specific gates by substitution. In particular, second-order defects of a direct difference approximation to the target bracket are not a finite BFV nilpotency theorem.

9 A real Lorentzian parent and its classical master equation

9.1 Fields, action, and polynomial structure

Fix $\eta_{IJ} = \mathrm{diag}(-1,1,1,1)$ and $\epsilon_{0123} = 1$. The independent fields are a real coframe e^I, a real antisymmetric Lorentz connection ω^{IJ}, an Abelian connection a, an antisymmetric auxiliary scalar φ_{IJ}, and Dirac fields $\psi, \bar{\psi}$. The spinors are Grassmann odd. Define

$$R^{IJ} = \mathrm{d}\omega^{IJ} + \omega^I{}_K \wedge \omega^{KJ}, \qquad F = \mathrm{d}a, \tag{68}$$

$$\mathrm{vol}_e = \tfrac{1}{4!}\epsilon_{IJKL} e^I \wedge e^J \wedge e^K \wedge e^L, \qquad \Sigma_I = \tfrac{1}{3!}\epsilon_{IJKL} e^J \wedge e^K \wedge e^L. \tag{69}$$

Use $\{\gamma^I, \gamma^J\} = 2\eta^{IJ}$ and $\gamma_{IJ} = \frac{1}{2}[\gamma_I, \gamma_J]$. The covariant derivatives are

$$\mathcal{D}\psi = \mathrm{d}\psi + \tfrac{1}{4}\omega^{IJ}\gamma_{IJ}\psi - iqa\psi, \tag{70}$$

$$\mathcal{D}\bar{\psi} = \mathrm{d}\bar{\psi} - \tfrac{1}{4}\bar{\psi}\omega^{IJ}\gamma_{IJ} + iqa\bar{\psi}. \tag{71}$$

The parameter q here is electric charge. The determinant of a spatial metric below will be written $\det(q_{ab})$ to avoid confusing the two uses.

The parent action is $S_0 = S_{\mathrm{G}} + S_{\mathrm{M}}^{(1)} + S_{\mathrm{D}}$, with

$$S_{\mathrm{G}} = \frac{1}{4\kappa}\int \epsilon_{IJKL} e^I \wedge e^J \wedge R^{KL} - \frac{\Lambda}{\kappa}\int \mathrm{vol}_e, \tag{72}$$

$$S_{\mathrm{M}}^{(1)} = \int \left(\tfrac{1}{2}\varphi_{IJ} F \wedge e^I \wedge e^J - \tfrac{1}{4}\varphi_{IJ}\varphi^{IJ}\,\mathrm{vol}_e\right), \tag{73}$$

$$S_{\mathrm{D}} = \frac{i}{2}\int \Sigma_I \wedge \left(\bar{\psi}\gamma^I \mathcal{D}\psi - (\mathcal{D}\bar{\psi})\gamma^I \psi\right) - m\int \mathrm{vol}_e\, \bar{\psi}\psi. \tag{74}$$

This formulation uses no inverse tetrad in its defining action. Its largest monomials have six field factors. Nondegeneracy becomes necessary when the auxiliary variables are eliminated or a metric interpretation is imposed.

The geometric fields are real, but the physical spinor convention needs a separate check. The displayed action retains the inherited Dirac-type polynomial gauge model. With the stated Clifford sign and kinetic factor i, the flat-space operators $i\gamma^I\partial_I \pm m$ have product

$$(i\gamma^I\partial_I + m)(i\gamma^J\partial_J - m) = -\Box - m^2, \qquad \Box = \eta^{IJ}\partial_I\partial_J. \tag{75}$$

Thus its squared field equation is $\Box + m^2 = 0$, whereas the conventional real massive dispersion in this mostly-plus metric uses $\Box - m^2 = 0$. A consistent translation of the Clifford, kinetic, mass, and adjoint conventions is required before identifying this sector with conventional physical massive Dirac matter. Gauge invariance and the graded master-equation argument do not require that physical identification, and they do not resolve it. The massless algebraic sector does not have this squared-mass discrepancy.

Variation of φ gives the unambiguous wedge equation

$$F \wedge e^I \wedge e^J = \varphi^{IJ}\,\mathrm{vol}_e\,. \tag{76}$$

Under the standard Lorentzian Hodge convention $\alpha \wedge \star\beta = \langle \alpha, \beta\rangle\,\mathrm{vol}_e$, this equation identifies φ with minus the internal components of $\star F$. For example, $F = Ee^0 \wedge e^1$ has $\star F = -Ee^2 \wedge e^3$, whereas the

wedge equation gives $\varphi^{23} = E$. The inherited shorthand $\varphi = \star F$ needs this sign translation. Substitution of the wedge equation into the action gives the conventional result

$$S^{(1)}_{\mathrm{M}}\big|_{\varphi} = -\tfrac{1}{4}\int F_{IJ}F^{IJ}\,\mathrm{vol}_e\,. \tag{77}$$

The substitution repairs the Hodge shorthand while retaining the on-shell Maxwell action.

9.2 Gauge transformations and graded conventions

For even parameters (ξ, λ, u), the coframe transforms by its Lie derivative and Lorentz rotation, the connection by $\mathcal{L}_\xi\omega - D_\omega\lambda$, and the Abelian potential by $\mathcal{L}_\xi a + \mathrm{d}u$. The spinor transforms by $\mathcal{L}_\xi\psi + \frac{1}{4}\lambda^{IJ}\gamma_{IJ}\psi + iqu\psi$, with the dual action on $\bar\psi$. All internal contractions in (74) are invariant. A diffeomorphism changes its four-form density by a Lie derivative, which integrates to zero on a periodic domain or under stated boundary conditions.

The parameter algebra combines vector fields with Lorentz and Abelian transformations in a semidirect product. In the inherited field-space commutator convention,

$$[\delta_U, \delta_V]_{\mathrm{fld}} = -\delta_{[U,V]_{\mathrm{par}}}, \tag{78}$$

where

$$[(\xi,\lambda,u),(\eta,\mu,v)]_{\mathrm{par}} = \big([\xi,\eta], \mathcal{L}_\xi\mu - \mathcal{L}_\eta\lambda + [\lambda,\mu], \mathcal{L}_\xi v - \mathcal{L}_\eta u\big)\,. \tag{79}$$

The minus sign is a convention for these field-space generators. It fixes the plus sign in the ghost differential below.

Let c, χ, υ be odd ghosts and let s be a left odd derivation. All ghosts multiplying component fields are placed on the left. Set

$$sc^\mu = c^\nu\partial_\nu c^\mu, \qquad s\chi = \mathcal{L}_c\chi + \chi^2, \qquad s\upsilon = \mathcal{L}_c\upsilon. \tag{80}$$

On the ordinary fields use the gauge variations with ghosts in place of even parameters, interpreted with this left-placement rule. In particular,

$$s\psi = \mathcal{L}_c\psi + \tfrac{1}{4}\chi^{IJ}\gamma_{IJ}\psi + iq\upsilon\psi, \tag{81}$$

$$s\bar\psi = \mathcal{L}_c\bar\psi - \tfrac{1}{4}\chi^{IJ}\bar\psi\gamma_{IJ} - iq\upsilon\bar\psi. \tag{82}$$

The ordering in the second line matters: both χ and $\bar\psi$ are odd. The expression $-\bar\psi\chi$ is not the same expression. Moving a ghost through an odd dual field changes its sign. The inherited formula with right-placed ghosts is not used without this correction.

Nilpotency follows from the graded Jacobi identity for the parameter algebra and its action on each representation. A simple internal check makes the sign visible. Write the Lorentz representation-valued ghost as $g = \frac{1}{4}\chi^{IJ}\gamma_{IJ}$ and suppress diffeomorphisms and the Abelian factor. Then $sg = g^2$, $s\psi = g\psi$, and the dual rule is $s\bar\psi = \bar\psi g$ when the odd ghost has been moved to the right. Thus

$$s^2\psi = g^2\psi - g(g\psi) = 0, \qquad s^2\bar\psi = (\bar\psi g)g - \bar\psi g^2 = 0. \tag{83}$$

Using $s\bar\psi = -\bar\psi g$ with the same left-derivation convention instead gives $2\bar\psi g^2$. A commuting-parameter covariance test does not detect this Grassmann sign error.

9.3 The classical BV action

For all fields and ghosts Ξ, introduce antifields $\Xi^\ddagger$ with opposite parity and ghost number $-1-\mathrm{gh}(\Xi)$. Use a canonical odd antibracket, with right and left derivatives and the indicated graded ordering. The minimal action can be written

$$S_{\mathrm{BV}} = S_0 + \int \sum_{\Xi} \langle \Xi^\ddagger, s\Xi \rangle. \tag{84}$$

Antisymmetric Lorentz indices are counted once, or equivalently their redundant sum is given its usual factor $1/2$. No higher-antifield terms are needed for this off-shell closed representation algebra.

Proposition 9.1 (Classical master equation)**.** *With the declared graded conventions,* $sS_0 = 0$ *and* $s^2 = 0$ *imply*

$$(S_{\mathrm{BV}}, S_{\mathrm{BV}}) = 0. \tag{85}$$

Proof. The bracket of S_0 with the antifield-linear terms is twice its gauge variation. The bracket among antifield-linear terms has coefficient $s^2\Xi$ at each antifield, including the ghost antifields. These coefficients vanish. There is no bracket contribution from S_0 with itself. This exhausts the antifield expansion. □

This is a classical construction. It contains no regulated BV Laplacian and establishes no quantum master equation. The odd-field ordering stated above is part of the construction, rather than a claim that the inherited numerical script already checked every Grassmann coefficient.

9.4 Resolved support and Lorentzian reality

Four-dimensional trigonometric polynomials obey the same support lemma as (48). A periodic four-dimensional grid is an algebraic regularization for these identities. It is not a claim that physical spacetime has periodic time or that it is globally hyperbolic. On a slab, the associated boundary terms must instead be retained.

For coefficients in $\mathcal{T}_K^{(4)}$ and a finite Grassmann algebra, the parent action has degree at most six, and the displayed antifield couplings have degree at most three. The master-equation and invariance coefficients have at most seven band-K factors. The sufficient rule $16K < N$ places them below Nyquist and transfers the classical identities coefficientwise. The rule concerns the polynomial parent. It does not automatically survive elimination of inverse tetrads or density square roots.

Reality is built into e and ω. For real Lorentz bivectors, the internal dual obeys $\star^2 = -\mathrm{id}$ and the projectors

$$P_\pm = \tfrac{1}{2}(\mathrm{id} \mp i\star) \tag{86}$$

satisfy $\overline{P_+ X} = P_- X$. Both chiral blocks are retained and related by conjugation. A real vector field with an existing unique flow preserves this real locus. This parent-level statement supplies a real classical starting point. It does not prove a missing finite reality-preserving step for the earlier isolated self-dual coordinates.

10 Boundary constraints and the additional moment-map gate

10.1 An exact time–space split

On $M = I \times \Sigma$ with orientation $\mathrm{d}t \wedge \mathrm{or}(\Sigma)$, write

$$e^I = e^I_\perp \mathrm{d}t + E^I, \qquad \omega^{IJ} = \omega^{IJ}_\perp \mathrm{d}t + \Omega^{IJ}, \qquad a = a_\perp \mathrm{d}t + A. \tag{87}$$

Let $\mathcal{R} = \mathrm{d}_\Sigma \Omega + \Omega \wedge \Omega$ and $\mathcal{F} = \mathrm{d}_\Sigma A$. The curvature and field strength split as

$$R = \mathrm{d}t \wedge (\dot{\Omega} - D_\Omega \omega_\perp) + \mathcal{R}, \qquad F = \mathrm{d}t \wedge (\dot{A} - \mathrm{d}_\Sigma a_\perp) + \mathcal{F}. \tag{88}$$

Define

$$B_{IJ} = \tfrac{1}{4\kappa} \epsilon_{KLIJ} E^K \wedge E^L, \tag{89}$$

$$\Pi = \tfrac{1}{2} \varphi_{IJ} E^I \wedge E^J, \qquad \sigma_I = \tfrac{1}{3!} \epsilon_{IJKL} E^J \wedge E^K \wedge E^L. \tag{90}$$

Direct exterior multiplication gives the spatial symplectic potential

$$\begin{aligned} \Theta_\Sigma = \int_\Sigma [& B_{IJ} \wedge \delta \Omega^{IJ} + \Pi \wedge \delta A \\ & - \tfrac{i}{2} \sigma_I (\bar{\psi} \gamma^I \delta \psi - \delta \bar{\psi} \gamma^I \psi)]. \end{aligned} \tag{91}$$

The minus sign in the spinor term is fixed by the time–space orientation and by moving $\mathrm{d}t$ past the spatial three-form. It must be retained when comparing the spacetime action and its canonical chart.

The temporal coframe multiplies the Einstein–matter three-forms $\mathcal{C}_I = \mathcal{C}^{\mathrm{G}}_I + \mathcal{C}^{\mathrm{M}}_I + \mathcal{C}^{\mathrm{D}}_I$, where

$$\mathcal{C}^{\mathrm{G}}_I = \tfrac{1}{2\kappa} \epsilon_{IJKL} E^J \wedge \mathcal{R}^{KL} - \tfrac{\Lambda}{\kappa} \sigma_I, \tag{92}$$

$$\mathcal{C}^{\mathrm{M}}_I = \varphi_{IJ} \mathcal{F} \wedge E^J - \tfrac{1}{4} \varphi_{KL} \varphi^{KL} \sigma_I, \tag{93}$$

$$\mathcal{C}^{\mathrm{D}}_I = \tfrac{i}{4} \epsilon_{JIKL} E^K \wedge E^L \wedge (\bar{\psi} \gamma^J \mathcal{D}_\Sigma \psi - (\mathcal{D}_\Sigma \bar{\psi}) \gamma^J \psi) - m \sigma_I \bar{\psi} \psi. \tag{94}$$

Spatial integration by parts identifies the Lorentz and electric Gauss forms

$$\mathcal{G}_{IJ} = D_\Omega B_{IJ} - \tfrac{i}{8} \sigma_K \bar{\psi} \{\gamma^K, \gamma_{IJ}\} \psi, \tag{95}$$

$$\mathcal{G}_{\mathrm{U}} = \mathrm{d}_\Sigma \Pi - q \sigma_I \bar{\psi} \gamma^I \psi. \tag{96}$$

Consequently the multiplier equations are $\mathcal{C}_I = \mathcal{G}_{IJ} = \mathcal{G}_{\mathrm{U}} = 0$. These are exact formulas for the split action. The auxiliary constitutive equations and the simplicity relations are additional primary relations; they are not extra first-class gauge generators.

10.2 What is explicit in the internal BFV sector

On a regular symplectic reduction where internal gauge transformations have the stated moment maps, their algebra is Lorentz plus Abelian:

$$\{G_{\mathrm{L}}[\lambda], G_{\mathrm{L}}[\mu]\} = G_{\mathrm{L}}[[\lambda, \mu]], \qquad \{G_{\mathrm{L}}, G_{\mathrm{U}}\} = \{G_{\mathrm{U}}, G_{\mathrm{U}}\} = 0. \tag{97}$$

Its canonical BFV charge is explicit:

$$\Omega_{\text{int}} = G_{\text{L}}[\chi] + G_{\text{U}}[\upsilon] - \tfrac{1}{2}\langle \pi_\chi, [\chi, \chi]\rangle. \tag{98}$$

The moment-map relation cancels the terms quadratic in ghosts, and the Lorentz Jacobi identity cancels the cubic terms. Thus $\{\Omega_{\text{int}}, \Omega_{\text{int}}\} = 0$ under that regular moment-map hypothesis. This finite Lie-algebra mechanism is more explicit than a claim about a complete gravitational boundary charge.

10.3 Normal deformations require more boundary data

For a spatial vector field η, the Cartan identities give

$$\mathcal{L}_\eta \Omega = \iota_\eta \mathcal{R} + D_\Omega(\iota_\eta \Omega), \qquad \mathcal{L}_\eta A = \iota_\eta \mathcal{F} + \mathrm{d}_\Sigma(\iota_\eta A). \tag{99}$$

Accordingly, a covariant tangential generator acquires Lorentz and Abelian compensators to generate an ordinary spatial Lie derivative. The source records the expression

$$J_\Sigma[c, \chi, u] = C[\zeta(c)] + G_{\text{L}}[\chi - \iota_c \Omega] + G_{\text{U}}[u + \iota_c A]. \tag{100}$$

For tangential c, $\zeta(c) = \iota_c E$ is a spatial contraction. For a normal deformation, a spatial one-form does not by itself supply that contraction. One must specify the embedding variables and conjugate momenta, their symplectic extension, and the relevant normal field data. One must then quotient the presymplectic kernel and prove that the gauge action and its Hamiltonian descend.

Proposition 10.1 (Conditional boundary BFV construction)**.** *Suppose a regular derived boundary symplectic space, including the required embedding data, has been constructed. Suppose its full normal and tangential gauge action has an equivariant moment map J_Σ for the declared parameter algebra. Add canonical ghosts and their momenta, and let the ghost Hamiltonian generate the Chevalley–Eilenberg differential in the same conventions. Then their sum Ω_Σ obeys*

$$\{\Omega_\Sigma, \Omega_\Sigma\} = 0. \tag{101}$$

Proof. The bracket of the two moment-map terms reproduces the parameter bracket. The cross bracket with the ghost Hamiltonian cancels that term. The bracket of the ghost Hamiltonian with itself vanishes by the graded Jacobi identity. Equivariance excludes a central defect in the moment-map relation. □

This proposition states the exact missing implication. Off-shell closure of the bulk BRST transformations is insufficient to prove its boundary moment-map hypotheses. The action-split computations establish the split and tangential Cartan identities. They do not construct every normal embedding variable or verify a complete boundary symplectic reduction. The boundary result remains conditional on the hypotheses printed in its statement.

10.4 ADM ghosts and spectral headroom

On a nondegenerate spacelike slice, let $q_{ab} = \eta_{IJ} E^I_a E^J_b$ and define the densitized normal

$$\tilde{n}_I = \tfrac{1}{3!}\epsilon_{IJKL}\epsilon^{abc} E^J_a E^K_b E^L_c. \tag{102}$$

It satisfies $\tilde{n}_I E^I_a = 0$ and $-\eta^{IJ}\tilde{n}_I\tilde{n}_J = \det(q_{ab})$. A parameter $\zeta^I = \nu\tilde{n}^I + \eta^a E^I_a$ yields a densitized lapse ν and shift η. On a correctly reduced boundary space the expected normal bracket has structure shift

$$S^a(\nu,\mu) = \det(q_{cd})q^{ab}(\nu\partial_b\mu - \mu\partial_b\nu). \tag{103}$$

The improved spatial generator differs from the covariant vector generator by its internal compensators. This is the usual reason that the normal bracket may contain both spatial and Gauss terms in the first-order chart.

The map from covariant ghosts to lapse, shift, and internal ghosts depends on E, Ω, A. A canonical lift must transform ghost momenta and, generally, physical momenta as well. It may generate higher ghost terms. Merely replacing the ghosts in a quadratic formula does not prove that the new formula is canonical or nilpotent. The complete charge must be carried through that lift, rather than inferred from its ghost-linear coefficients.

The source estimates at most nine band-K factors in its polynomial ADM representation, so at most sixteen in a bracket. Its sufficient rule is $32K < N$. This explains the additional spectral room required by the densitized normal. It does not discharge the boundary construction, the canonical ghost lift, or rational inverses created after reduction. Those prerequisites are independent of a degree count.

11 A regular real canonical chart

11.1 Time gauge and simplicity pullback

Choose an oriented time-gauge patch

$$E^0 = 0, \qquad E^i = e_a{}^i\,\mathrm{d}x^a, \qquad e := \det(e_a{}^i) > 0. \tag{104}$$

The positive determinant fixes the orientation and ensures invertibility. Define

$$\widetilde{E}^a{}_i = \tfrac{1}{2}\epsilon^{abc}\epsilon_{ijk}e_b{}^j e_c{}^k, \qquad \widetilde{E}^a{}_i e_a{}^j = e\delta_i{}^j. \tag{105}$$

Let $K^i = \Omega^{0i}$ and $\Gamma^i = -\frac{1}{2}\epsilon^i{}_{jk}\Omega^{jk}$. The simplicity pullback gives

$$B_{ij} = 0, \qquad B_{0i} = \tfrac{1}{4\kappa}\epsilon_{ijk}E^j \wedge E^k, \qquad \int_\Sigma B_{IJ}\wedge\delta\Omega^{IJ} = \frac{1}{\kappa}\int_\Sigma \widetilde{E}^a{}_i\delta K_a{}^i\,\mathrm{d}^3x. \tag{106}$$

The antisymmetric IJ sum includes both $0i$ and $i0$, accounting for the factor in the last identity. There is no $\delta\Gamma$ term. Those directions lie in the kernel of the pulled-back gravitational presymplectic form and are fixed by the algebraic connection equation on the regular sector.

Similarly, with $\mathcal{E}^a = \frac{1}{2}\epsilon^{abc}\Pi_{bc}$,

$$\int_\Sigma \Pi\wedge\delta A = \int_\Sigma \mathcal{E}^a\delta A_a\,\mathrm{d}^3x. \tag{107}$$

The Maxwell auxiliary equation then gives its metric Hamiltonian

$$\mathcal{H}_{\mathrm{M}} = \frac{1}{2e}q_{ab}(\mathcal{E}^a\mathcal{E}^b + \mathcal{B}^a\mathcal{B}^b), \qquad \mathcal{B}^a = \tfrac{1}{2}\epsilon^{abc}\mathcal{F}_{bc}. \tag{108}$$

The inverse $1/e$ is well defined in the patch and singular at its boundary.

11.2 Fermionic primary constraints and half densities

Impose the inherited algebraic spinor relation $\bar{\psi} = \psi^\dagger\gamma^0$, where $(\gamma^0)^2 = -\mathrm{id}$. Its phase convention must be included in the physical translation just described. In time gauge $\sigma_0 = e\,\mathrm{d}^3x$ and $\sigma_i = 0$. Equation (91) becomes

$$\Theta_{\mathrm{D}} = \frac{i}{2}\int_\Sigma e(\psi^\dagger\delta\psi - \delta\psi^\dagger\psi)\,\mathrm{d}^3x. \tag{109}$$

In the inherited right–left graded canonical convention, the primary matrix is

$$\{\chi_A(x), \bar{\chi}_B(y)\}_{\mathrm{P}} = -ie(x)\delta_{AB}\delta^3(x-y). \tag{110}$$

It is invertible exactly where $e \neq 0$. Set $\Psi = e^{1/2}\psi$ and $\Psi^\dagger = e^{1/2}\psi^\dagger$. Differentiation gives

$$\Psi^\dagger\delta\Psi = e\psi^\dagger\delta\psi + \tfrac{1}{2}\psi^\dagger\psi\,\delta e, \tag{111}$$

$$\delta\Psi^\dagger\Psi = e\delta\psi^\dagger\psi + \tfrac{1}{2}\psi^\dagger\psi\,\delta e. \tag{112}$$

The determinant terms cancel, so

$$\Theta_{\mathrm{D}} = \frac{i}{2}\int_\Sigma(\Psi^\dagger\delta\Psi - \delta\Psi^\dagger\Psi)\,\mathrm{d}^3x, \qquad \{\Psi_A(x), \Psi^\dagger_B(y)\}_{\mathrm{D}} = -i\delta_{AB}\delta^3(x-y). \tag{113}$$

The half density removes the field-dependent coefficient in the fermionic primary bracket without shifting the gravitational momentum. It requires the chosen real positive square root of e, so it is a local coordinate choice on (104).

11.3 Algebraic torsion and the reduced phase space

For the symmetric Dirac action, define $\gamma_5 = i\gamma^0\gamma^1\gamma^2\gamma^3$ and $J^I_5 = \bar{\psi}\gamma^I\gamma_5\psi$. In the parent convention, the connection equation has the algebraic solution

$$T^I = \tfrac{\kappa}{4}\epsilon^I{}_{JKL}J^J_5 e^K \wedge e^L, \qquad C^{IJ} = \tfrac{\kappa}{4}\epsilon^{IJ}{}_{KL}J^K_5 e^L, \tag{114}$$

with $\omega = \mathring{\omega}(e) + C$. The solution is unique on the nondegenerate coframe sector. It fixes

$$\Gamma_a{}^i = \mathring{\Gamma}_a{}^i(E) - \tfrac{\kappa}{4}J^0_5 e_a{}^i. \tag{115}$$

One may retain torsion in K and in the substituted connection, or use a torsion-free representation with the induced local four-spinor interaction. In the declared action convention the latter action contribution is $-\frac{3\kappa}{16}\int J_{5I}J^I_5\,\mathrm{vol}_e$. Its sign belongs to that convention; changing the Clifford, orientation, or Dirac-action conventions requires translating it.

The resulting real canonical superphase space has coordinates

$$(K_a{}^i, \widetilde{E}^a{}_i), \qquad (A_a, \mathcal{E}^a), \qquad (\Psi, \Psi^\dagger), \tag{116}$$

and one-form

$$\Theta_{\mathrm{red}} = \int_\Sigma \mathrm{d}^3x\,\left[\frac{1}{\kappa}\widetilde{E}^a{}_i\delta K_a{}^i + \mathcal{E}^a\delta A_a + \frac{i}{2}(\Psi^\dagger\delta\Psi - \delta\Psi^\dagger\Psi)\right]. \tag{117}$$

Its even canonical dimension is 24 per spatial point: eighteen gravitational and six electromagnetic coordinates. The odd primary block has eight components in the stated complexified graded presen-

tation, with the real spinor relation imposed. These counts are local phase-space counts, not a count of globally gauge-quotiented physical degrees of freedom. Rotation, electric, spatial, and Hamiltonian constraints are still first class.

For example, the residual electric constraint is

$$\mathcal{G}_{\mathrm{U}}^{\mathrm{red}} = \partial_a \mathcal{E}^a + q\Psi^\dagger\Psi, \qquad Q[u] = -\int_\Sigma u\mathcal{G}_{\mathrm{U}}^{\mathrm{red}}\,\mathrm{d}^3x. \tag{118}$$

It generates $\delta_u A_a = \partial_a u$ and $\delta_u \Psi = iqu\Psi$. The half-density Lie derivative in the spatial generator is

$$\mathcal{L}_\eta^{(1/2)}\Psi = \eta^a\partial_a\Psi + \tfrac{1}{2}(\partial_a\eta^a)\Psi. \tag{119}$$

This term is forced by the density weight. Treating Ψ as a scalar would give a different moment map.

11.4 Descent and the boundary of the chart

Let ι denote the local solution map for the second-class relations and let $\{\cdot,\cdot\}_{\mathrm{D}}$ be its regular Dirac bracket. If an already constructed boundary BFV vector field is tangent to that second-class surface, and its charge descends through the regular reduction, then

$$\Omega_{\mathrm{red}} = \iota^*\Omega_\Sigma, \qquad \{\Omega_{\mathrm{red}}, \Omega_{\mathrm{red}}\}_{\mathrm{D}} = 0. \tag{120}$$

The implication follows from compatibility of the reduced bracket with the Dirac bracket and the vanishing fine master equation. It does not construct the missing full boundary charge. Fixing time gauge also removes boost gauge directions, so the tangent action must be the residual action or a correctly compensated one.

Degenerate triads, changes of orientation, and singular gauge patches are outside this chart. At $e = 0$ the primary fermionic matrix collapses and the coordinate $e^{1/2}\psi$ ceases to be an invertible change of variables. This is a stratified boundary of the reduction, not a small numerical error. A global reduction would need compatible charts and a treatment of these singular strata.

There is also a spectral boundary. Even when ${e_a}^i$ are trigonometric polynomials, e^{-1}, $e^{1/2}$, and the torsion-free connection are generally not finite Fourier polynomials. The unreduced degree rules $16K < N$ and $32K < N$ therefore cannot simply be reapplied to the reduced chart. One may evaluate its finite-dimensional regular functions directly, or retain auxiliary polynomial variables, but either choice needs its own approximation and support analysis.

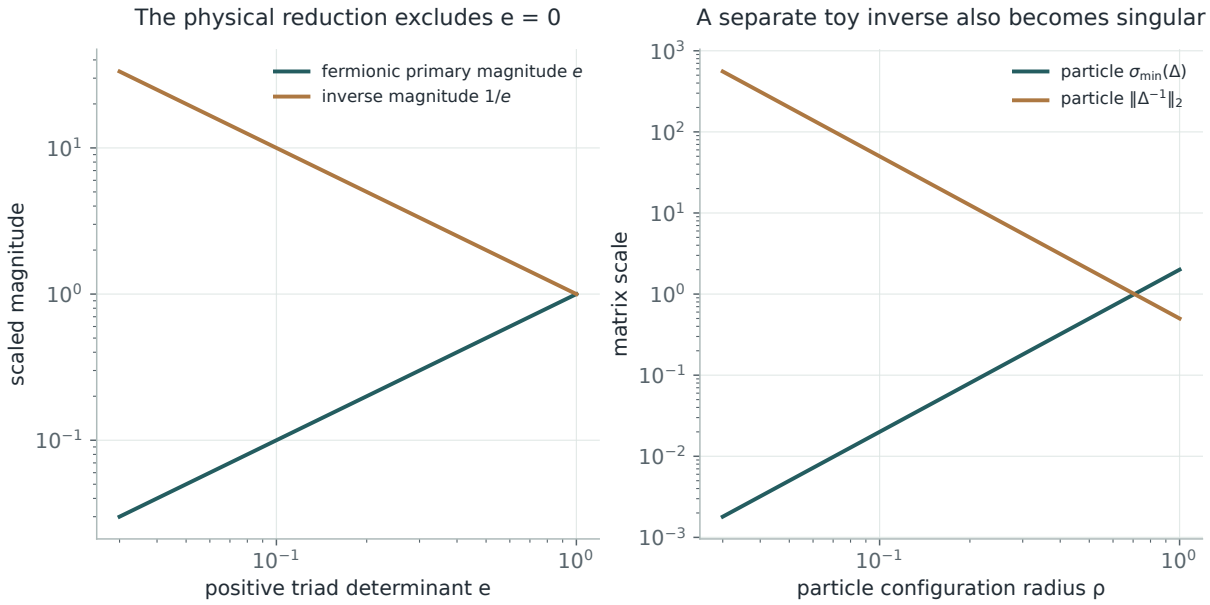


Figure 3. The regular triad sector and a separate finite particle gauge patch. Invertibility is a hypothesis of each reduction, and a singular boundary changes the rank of its constraint matrix.

12 Exactly preserved steps on a finite regular gauge patch

12.1 Endpoint Dirac projection

Consider a finite-dimensional real canonical chart $z \in \mathbb{R}^{2n}$ with constant skew matrix J. Let $\Phi = (\Phi_1, \dots, \Phi_{2m})$ collect independent first-class constraints and regular gauge conditions, or more generally a regular second-class set. Assume these are smooth on a neighborhood of the whole segment between the two step endpoints. For a smooth autonomous function F, define the segment discrete gradient

$$\overline{\nabla} F(z_0, z_1) = \int_0^1 \nabla F((1-s)z_0 + s z_1)\, \mathrm{d}s. \tag{121}$$

The ordinary fundamental theorem of calculus gives

$$F(z_1) - F(z_0) = \overline{\nabla} F(z_0, z_1)^T (z_1 - z_0). \tag{122}$$

Let G have rows $\overline{\nabla}\Phi_j^T$. The endpoint constraint matrix is $\Delta_\Delta = GJG^T$. Require this matrix to be invertible and define

$$J_{\mathrm{D}}^{\Delta} = J - JG^T \Delta_\Delta^{-1} GJ. \tag{123}$$

Direct multiplication proves $GJ_{\mathrm{D}}^{\Delta} = 0$. Since Δ_Δ and its inverse are skew, transposition proves $(J_{\mathrm{D}}^{\Delta})^T = -J_{\mathrm{D}}^{\Delta}$. The endpoint tensor is a projection of the canonical direction into the kernel of all discrete constraint gradients. It need not define a global Poisson bracket or satisfy a Jacobi identity away from the step. Those properties are unnecessary for the preservation argument.

Theorem 12.1 (Exact finite-step preservation)**.** *Let H be a smooth autonomous Hamiltonian. Every exact solution of*

$$z_1 - z_0 = \Delta t\, J_{\mathrm{D}}^{\Delta}(z_0, z_1)\overline{\nabla} H(z_0, z_1) \tag{124}$$

with invertible Δ_{Δ} satisfies

$$\Phi(z_1) = \Phi(z_0), \qquad H(z_1) = H(z_0). \tag{125}$$

Proof. The componentwise chain rule gives $\Phi(z_1) - \Phi(z_0) = G(z_1 - z_0)$, which vanishes because $GJ_{\mathrm{D}}^{\Delta} = 0$. For the Hamiltonian,

$$H(z_1) - H(z_0) = \Delta t\, \overline{\nabla} H^T J_{\mathrm{D}}^{\Delta} \overline{\nabla} H = 0 \tag{126}$$

by skewness. □

The theorem preserves the initial values of the constraints. It does not repair inconsistent initial data. For explicitly time-dependent gauges or $H(z, t)$, ordinary energy is not a conserved quantity. One must formulate the desired balance law or enlarge the phase space by a clock and its conjugate momentum before applying an autonomous preservation theorem.

12.2 A worked particle constraint

In $\mathbb{R}^6$ with $z = (q, p)$ and canonical J, set

$$\Phi_1 = q \cdot q - 1, \qquad \Phi_2 = q \cdot p. \tag{127}$$

At one point the ordinary Dirac matrix is

$$\Delta(z) = \begin{pmatrix} 0 & 2\|q\|^2 \\ -2\|q\|^2 & 0 \end{pmatrix}. \tag{128}$$

It is invertible whenever $q \neq 0$. On the intended constraint surface, q lies on the unit sphere and p is tangent to that sphere. Writing $\bar{q} = (q_0 + q_1)/2$ and $\bar{p} = (p_0 + p_1)/2$, the exact discrete gradients are

$$\overline{\nabla}\Phi_1 = (2\bar{q}, 0), \qquad \overline{\nabla}\Phi_2 = (\bar{p}, \bar{q}). \tag{129}$$

Consequently Δ_{Δ} has the same displayed form with q replaced by $\bar{q}$. Endpoints on the unit sphere can have $\bar{q} = 0$ if they are antipodal. Endpoint admissibility therefore does not imply endpoint-matrix invertibility for an arbitrary large step. Small local steps avoid this singular pair under the regularity hypotheses below.

For $H(q, p) = \frac{1}{2}\|p\|^2 + V(q)$, the segment formula gives an exact discrete gradient of the potential. Polynomial V makes the segment integral a finite algebraic expression. The resulting implicit equations preserve both sphere constraints and H when solved exactly. This is a genuine finite constrained mechanical model. It is a diagnostic for the step mechanism, not an Einstein–Maxwell–Dirac simulation.

12.3 Solver residuals have a precise effect

For an approximate endpoint, define its equation residual by

$$r = (z_1 - z_0) - \Delta t\, J_{\mathrm{D}}^{\Delta}\overline{\nabla} H. \tag{130}$$

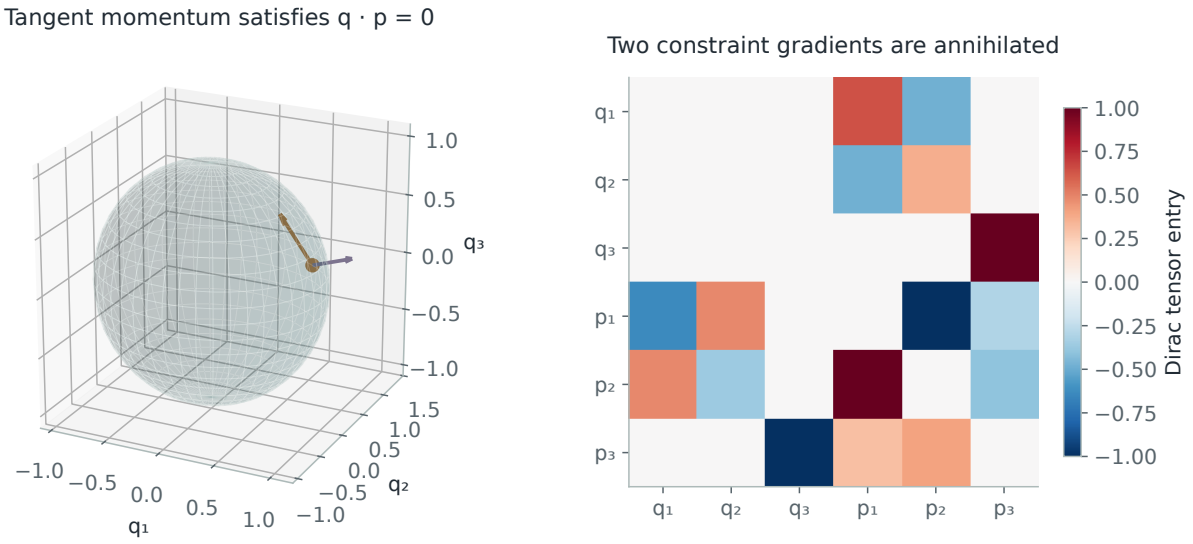


Figure 4. A three-dimensional sphere and its tangent directions, together with the finite Dirac projection. The projection annihilates the specified constraint gradients while retaining physical tangent directions.

If the discrete gradients and tensor are evaluated exactly, the same proof gives

$$\Phi(z_1) - \Phi(z_0) = Gr, \qquad H(z_1) - H(z_0) = \overline{\nabla} H^T r. \tag{131}$$

Thus $\|\Phi(z_1)-\Phi(z_0)\| \le \|G\|\,\|r\|$ and $|H(z_1)-H(z_0)| \le \|\overline{\nabla}H\|\,\|r\|$. These estimates explain why an incompletely solved step need not preserve the invariants. Numerical quadrature of nonpolynomial segment integrals and floating-point linear algebra contribute additional errors and must also be measured. A printed solver tolerance is not the same quantity as the actual equation residual.

The particle certificate records 1000 steps at $\Delta t = 0.1$, with energy error about 7.8×10^{-16}, constraint error about 3.9×10^{-15}, and equation residual about 1.1×10^{-16}. A Runge–Kutta negative control on the continuous Dirac vector field has errors around 10^{-6}. Those values concern that fixed test. They do not establish exact floating-point preservation or a uniform gravity stability theorem. The accompanying atlas reruns a declared finite particle example and also tests incomplete implicit solves.

12.4 Local existence, boundedness, and refinement

At $\Delta t = 0$ and $z_1 = z_0$, the endpoint tensor reduces to the ordinary Dirac tensor if the point constraint matrix is invertible. The derivative with respect to z_1 of the left side of (124) minus its right side is then the identity. If the segment data are smooth in a neighborhood and the inverse matrix stays bounded, the implicit-function theorem gives a unique nearby step for sufficiently small Δt. It does not prove existence of a unique endpoint at arbitrary step size or after leaving that neighborhood.

If the gauge-fixed autonomous Hamiltonian satisfies

$$H(z) \ge a\|z\|^p - b, \qquad a > 0, \quad p > 0, \tag{132}$$

exact preservation yields

$$\|z_n\| \le \left(\frac{H(z_0) + b}{a} \right)^{1/p} \tag{133}$$

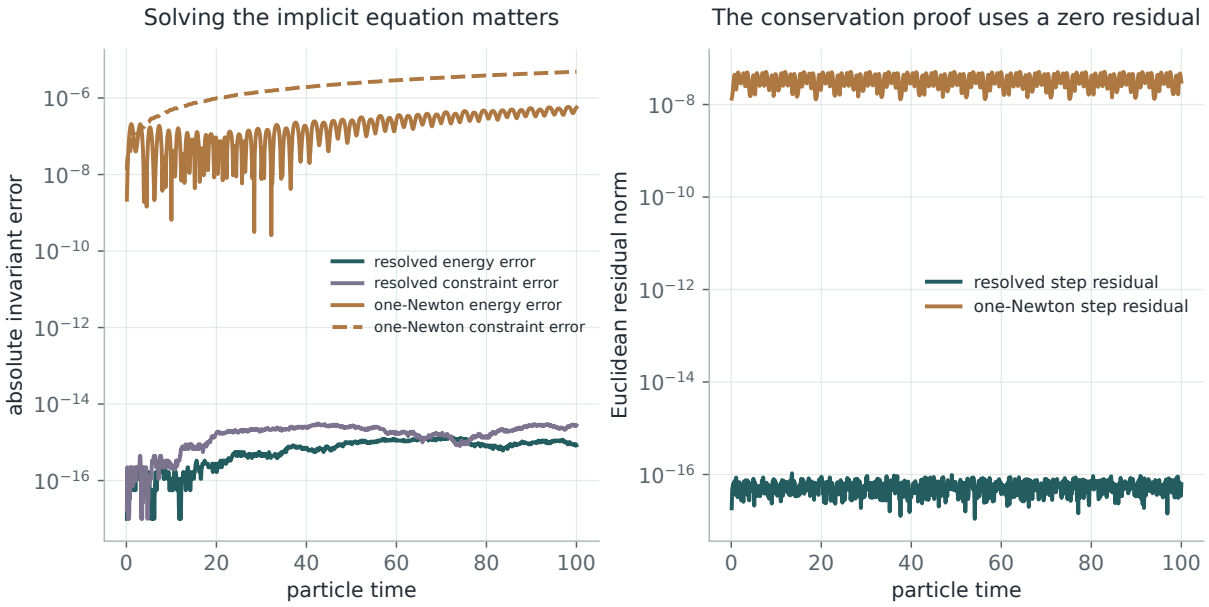


Figure 5. A fresh finite particle experiment contrasts resolved implicit steps with an incomplete-solve control. Its numerical residuals test the preservation mechanism within the specified model and range.

for every continued step in the regular patch. This coercivity is a separate assumption. It is not automatic for a gravitational constraint or a Hamiltonian that vanishes on its constraint surface. Bounded coordinates do not ensure that the gauge matrix stays nonsingular or that the triad determinant stays away from zero.

For an embedding $I_{K\to L}$ and a level-L step U_L, the adaptive map is $U_L \circ I_{K\to L}$. The level-L step preserves the invariants of its embedded initial state. To preserve the intended old data, the embedding must also satisfy the appropriate pullback relations for the Hamiltonian, constraints, gauge conditions, and symplectic form. A generic zero padding or a generic change of gauge does not prove all those relations.

A uniform refinement theorem additionally requires level-independent derivative bounds, inverse Dirac-matrix bounds, coercivity, and a consistency estimate for the actual dynamics. The source leaves those estimates open for the full gravity–matter tower. Strict finite-range implementation is another independent problem: a Fourier derivative couples all grid sites, even though its algebraic operations can be exactly resolved.

13 Elimination, measure balance, and the remaining quantum gate

13.1 A nonlinear Gaussian reduction in closed form

The small elimination model in the transfer chapter explains why an effective action and an effective measure must be transported together. Let $m, n > 0$ and

$$S_F(x_1, x_2, y, z) = \tfrac{1}{2}(x_2 - x_1 y)^2 + \tfrac{m}{2}(y - z)^2 + \tfrac{n}{2} z^2. \tag{134}$$

The fine gauge vector $R_F = \partial_{x_1} + y\partial_{x_2}$ leaves S_F invariant and has zero ordinary divergence. The detail Hessian is

$$M_D = \begin{pmatrix} x_1^2 + m & -m \\ -m & m+n \end{pmatrix}, \qquad D = (m+n)x_1^2 + mn > 0. \tag{135}$$

Writing $\mu = mn/(m+n)$ and $d = x_1^2 + \mu$, the stationary detail values are $y_* = x_1x_2/d$ and $z_* = my_*/(m+n)$. Substitution gives

$$S_C = \frac{\mu x_2^2}{2d}, \qquad R_C = \partial_{x_1} + \frac{x_1x_2}{d}\partial_{x_2}. \tag{136}$$

Differentiation proves $R_C S_C = 0$. Gaussian integration, rather than stationary substitution alone, also gives the density $\rho_C \propto D^{-1/2}$, up to a constant. It satisfies

$$\operatorname{div} R_C = \frac{x_1}{d}, \qquad R_C \log \rho_C = -\frac{x_1}{d}. \tag{137}$$

Their sum vanishes. The effective measure therefore restores the divergence balance of the nonlinear induced gauge transformation.

Eliminating z first gives the effective coefficient μ for y^2 and the factor $(m+n)^{-1/2}$. Eliminating y next gives $d^{-1/2}$. Their product is $D^{-1/2}$, exactly as in direct elimination. For the Euclidean weight $e^{-S/\hbar}$, using a flat coarse measure instead records this determinant as $S_{C,\hbar} = S_C + \frac{\hbar}{2}\log D$, up to an additive constant. A nonlinear symmetry that preserves the classical action can still require a nonconstant induced measure.

13.2 The conditional BV pushforward statement

Suppose a regulated fine quantum BV theory already has compatible odd-symplectic and measure data and satisfies

$$\tfrac{1}{2}(S_F, S_F) - i\hbar\Delta_F S_F = 0. \tag{138}$$

Suppose the variables split into retained and detail sectors with $\Delta_F = \Delta_C + \Delta_D$, and choose an admissible detail Lagrangian $\mathcal{L}_D$. Define the coarse weight by

$$e^{iS_C/\hbar} = \mathcal{N}^{-1} \int_{\mathcal{L}_D} e^{iS_F/\hbar}. \tag{139}$$

When this integral is well defined in the declared analytic or perturbative sense and BV integration by parts has no boundary contribution,

$$\Delta_C e^{iS_C/\hbar} = -\mathcal{N}^{-1} \int_{\mathcal{L}_D} \Delta_D e^{iS_F/\hbar} = 0. \tag{140}$$

This is preservation of the quantum master equation under compatible elimination. It assumes the fine quantum master equation. The classical parent in (84) supplies none of the missing regulator, measure, or anomaly-control data required by (138).

Similarly, direct and iterated elimination agree when the integrals and product gauge fixings admit the relevant Fubini or formal pushforward composition. The quadratic special case is Schur elimination,

$$L_{\text{eff}} = L_{CC} - L_{CD}L_{DD}^{-1}L_{DC}, \tag{141}$$

with the required inverses. A semigroup of elimination maps under these hypotheses does not imply a physically local coarse action. Exact detail elimination can create long-range couplings.

Spherical strong fields

14 Mass coordinates and spherical electrovac

A strong gravitational field is a useful test of a constrained construction because coordinate singularities, physical curvature, and numerical constraint drift can occur at quite different places. The spherical models below keep those distinctions visible. The electrovac model supplies an exact curved spacetime whose coordinates move in time. The scalar model supplies dynamical matter and a numerically observed marginally trapped sphere. Their computational maps have different locality properties, and their evidence supports different conclusions.

The electrovac sector sets the Dirac field to zero and imposes spherical symmetry. It retains a real canonical Einstein–Maxwell system, but it has no nonspherical gravitational radiation. Work in units $G = c = 4\pi\epsilon_0 = 1$. A prime denotes a radial derivative and a dot denotes a time derivative. On a radial interval where $R > 0$ and $\Lambda > 0$, write

$$ds^2 = -N^2dt^2 + \Lambda^2(dr + N^r dt)^2 + R^2 d\Omega^2, \qquad \mathcal{A} = A_r dr + \Phi_e dt. \tag{142}$$

Here R is the areal radius, Λ determines proper radial distance, N is the lapse, and N^r is the radial shift. The Maxwell multiplier is denoted Φ_e to distinguish it from the scalar gradient used later. The bulk canonical action is

$$S = \int dt \int dr \left(P_\Lambda \dot{\Lambda} + P_R \dot{R} + Q\dot{A}_r - N\mathcal{H} - N^r\mathcal{D} - \Phi_e\mathcal{G}\right) + B_\partial. \tag{143}$$

The boundary term specifies the mass and electromagnetic boundary ensemble. Its treatment cannot be omitted when claiming a global Hamiltonian evolution, even though it does not change the following local bulk identity.

The gauge-improved radial momentum and electric constraints are

$$\mathcal{D} = P_R R' - \Lambda P'_\Lambda, \qquad \mathcal{G} = Q', \tag{144}$$

and the Hamiltonian constraint is

$$\mathcal{H} = -\frac{P_R P_\Lambda}{R} + \frac{\Lambda P_\Lambda^2}{2R^2} + \frac{RR''}{\Lambda} - \frac{RR'\Lambda'}{\Lambda^2} + \frac{R'^2}{2\Lambda} - \frac{\Lambda}{2} + \frac{\Lambda Q^2}{2R^2}. \tag{145}$$

The momenta follow the convention

$$P_\Lambda = -\frac{R}{N}(\dot{R} - N^r R'), \tag{146}$$

$$P_R = -\frac{1}{N}\left[\Lambda(\dot{R} - N^r R') + R(\dot{\Lambda} - (N^r\Lambda)')\right]. \tag{147}$$

These signs fix which Painlevé–Gullstrand branch describes inward future characteristics.

Definition 14.1 (Charged mass coordinate)**.** The charged Misner–Sharp coordinate in this canonical chart is

$$\mathcal{M} = \frac{R}{2}\left[1 - \left(\frac{R'}{\Lambda}\right)^2 + \left(\frac{P_\Lambda}{R}\right)^2\right] + \frac{Q^2}{2R}. \tag{148}$$

The charge contribution is included in $\mathcal{M}$ so that it is constant throughout an electrovac Reissner–Nordström solution.

Proposition 14.2 (Exact triangular mass identity)**.** *For differentiable canonical data in the regular chart,*

$$\mathcal{M}' = -\frac{R'}{\Lambda}\mathcal{H} - \frac{P_\Lambda}{R\Lambda}\mathcal{D} + \frac{Q}{R}\mathcal{G}. \tag{149}$$

Consequently, on a patch where $R'/\Lambda \neq 0$, the systems $(\mathcal{H}, \mathcal{D}, \mathcal{G}) = 0$ and $(\mathcal{M}', \mathcal{D}, \mathcal{G}) = 0$ have the same zero set.

Proof. Differentiating (148) gives

$$\begin{aligned}\mathcal{M}' = {} & \frac{R'}{2} - \frac{R'^3}{2\Lambda^2} - \frac{RR'R''}{\Lambda^2} + \frac{RR'^2\Lambda'}{\Lambda^3} + \frac{P_\Lambda P'_\Lambda}{R} - \frac{P_\Lambda^2 R'}{2R^2} \\ & + \frac{QQ'}{R} - \frac{Q^2R'}{2R^2}.\end{aligned} \tag{150}$$

Expanding the right side of (149) cancels the two P_R terms and produces precisely this expression. The constraint transformation is upper triangular with diagonal $(-R'/\Lambda, 1, 1)$, so it is invertible under the stated condition. □

This calculation changes constraint coordinates before discretization. It does not prove closure of an arbitrary pointwise discretization of the original hypersurface-deformation algebra. It also establishes equivalence of zero sets, rather than proving that every Poisson bracket of the replacement functions vanishes. The distinction matters when the reduced construction is compared with a general BFV system.

15 Horizon penetration and the regular radial chart

Fix $M > 0$, a real charge Q, and a branch $\varepsilon \in \{+1, -1\}$. On $2MR - Q^2 > 0$, define

$$p(R) = \varepsilon\sqrt{2MR - Q^2}, \qquad v(R) = \frac{p(R)}{R}. \tag{151}$$

The Painlevé–Gullstrand metric takes the form

$$ds^2 = -dt_P^2 + (dR + v(R)dt_P)^2 + R^2 d\Omega^2. \tag{152}$$

No inverse factor of $1 - 2M/R + Q^2/R^2$ appears. The metric can therefore cross a horizon while the lapse and radial spatial metric remain regular.

Take a smooth radial map $R = F(t, r)$ with $F_r > 0$ and identify $t_P = t$. Expanding $dR = F_t dt + F_r dr$ identifies

$$N = 1, \qquad \Lambda = F_r, \qquad N^r = \frac{F_t + v(F)}{F_r}. \tag{153}$$

The corresponding canonical data are

$$P_\Lambda = p(F), \qquad P_R = \frac{MF_r}{p(F)}, \qquad A_r = 0, \quad \Phi_e = -\frac{Q}{F}, \quad Q' = 0. \tag{154}$$

The sign in P_R is contained in $p(F)$.

Proposition 15.1 (A time-dependent exact gauge family)**.** *The data* (154) *satisfy the local electrovac constraints and canonical evolution equations wherever $F_r > 0$, $F > 0$, and $2MF - Q^2 > 0$.*

Proof. First, $R'/\Lambda = 1$ and $p(F)^2 = 2MF - Q^2$ imply $\mathcal{M} = M$. Since

$$P'_\Lambda = \frac{MF_r}{p(F)} = P_R, \tag{155}$$

the momentum constraint vanishes. The mass identity and $Q' = 0$ then give $\mathcal{H} = 0$.

For evolution, the radius equation gives

$$\dot{R} = -\frac{P_\Lambda}{R} + N^r R' = -v(F) + F_t + v(F) = F_t. \tag{156}$$

The radial-metric equation is

$$\dot{\Lambda} = \frac{\Lambda P_\Lambda}{R^2} - \frac{P_R}{R} + (N^r \Lambda)'. \tag{157}$$

Use $p_R = M/p$ and $v_R = M/(pR) - p/R^2$. The non-F_{tr} terms cancel, leaving $\dot{\Lambda} = F_{tr}$. The momentum equations in the $N = 1$ convention are

$$\dot{P}_\Lambda = N^r P'_\Lambda + \frac{1}{2}\left[1 - \frac{P_\Lambda^2 + Q^2}{R^2} - \left(\frac{R'}{\Lambda}\right)^2\right], \tag{158}$$

$$\dot{P}_R = \frac{\Lambda(P_\Lambda^2 + Q^2)}{R^3} - \frac{P_\Lambda P_R}{R^2} + (N^r P_R)' - \frac{R''}{\Lambda} + \frac{\Lambda' R'}{\Lambda^2}. \tag{159}$$

Substituting the data and differentiating $p(F)^2 = 2MF - Q^2$ reduces their right sides to

$$\frac{MF_t}{p(F)}, \qquad M\left[\frac{F_{tr}}{p(F)} - \frac{MF_r F_t}{p(F)^3}\right], \tag{160}$$

which are the direct time derivatives of the two momenta. Finally, $\Phi'_e = QF_r/F^2$ cancels $N\Lambda Q/R^2$ in $\dot{A}_r = N\Lambda Q/R^2 - \Phi'_e$. □

An arbitrary admissible F therefore creates coordinate motion through one fixed electrovac spacetime. Neither the mass nor the charge changes. A numerical reconstruction of this family is a nonlinear strong-field consistency test, but it is not matter collapse.

15.1 Charged horizons and the turning surface

Write

$$f(R) = 1 - \frac{2M}{R} + \frac{Q^2}{R^2}. \tag{161}$$

For $0 < |Q| < M$, its two positive roots are

$$r_{\pm} = M \pm \sqrt{M^2 - Q^2}. \tag{162}$$

At either root $p(R)^2 = R^2$, so the canonical quantities in (154) remain finite if F_r does. For $|Q| = M$, the two roots coincide at $R = M$. For $|Q| > M$ there is no real root. At $Q = 0$, the outer root is $2M$ and the formal inner root 0 is the singular endpoint, not a second regular horizon.

The square-root chart instead meets a turning surface at

$$R_{\text{turn}} = \frac{Q^2}{2M}. \tag{163}$$

There $p = 0$, and $P_R = MF_r/p$ becomes unbounded for nonzero F_r. This is a limitation of the chosen flat spatial slice and canonical chart. Horizon penetration does not remove the need for a positive distance from this turning surface. It also does not remove the physical curvature singularity at $R = 0$.

The invariant curvature diagnostic is

$$\mathcal{K} = \frac{48M^2}{R^6} - \frac{96MQ^2}{R^7} + \frac{56Q^4}{R^8}. \tag{164}$$

For $M = 1$, $Q = 0.8$, the radii are $r_+ = 1.6$, $r_- = 0.4$, and $R_{\text{turn}} = 0.32$. The benchmark interval $R \in [0.6, 6]$ crosses the outer horizon only. At its inner end $\mathcal{K} \simeq 199.6646853$, while $\mathcal{K}(r_+) \simeq 1.106262207$. High curvature and regular horizon-crossing coexist on this interval.

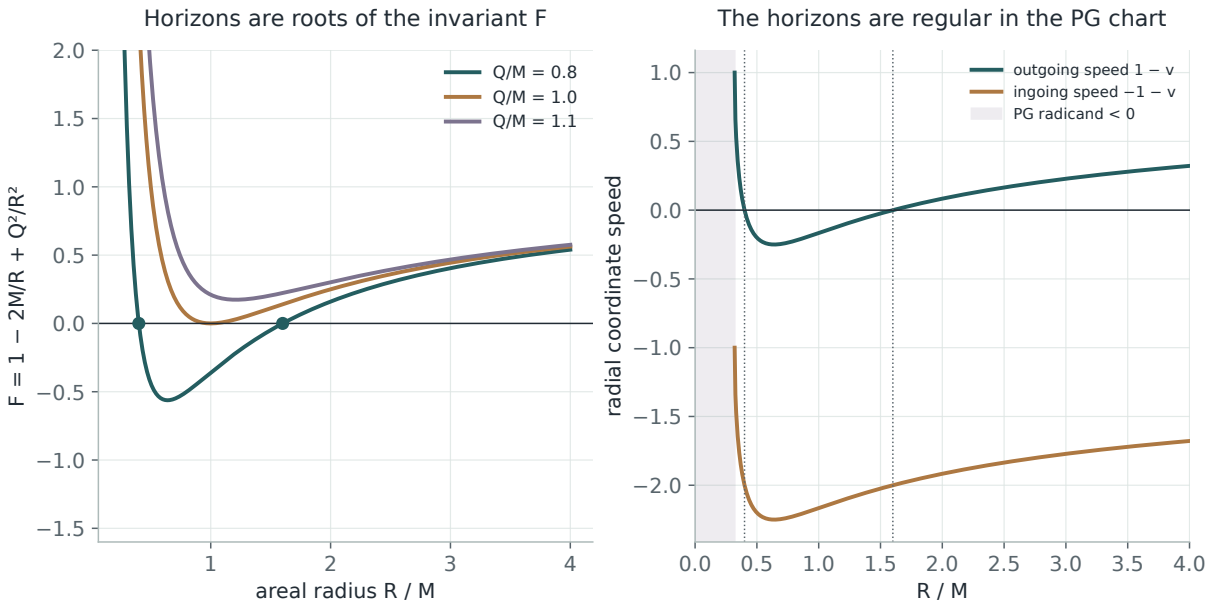


Figure 6. The horizon equation and the domain of the charged Painlevé–Gullstrand chart are different conditions. The benchmark crosses the outer horizon while remaining above the turning surface. The extremal and superextremal curves illustrate the change in the number and multiplicity of horizon roots.

15.2 Horizon conditioning

For a subextremal outer horizon, put $d = \sqrt{M^2 - Q^2} > 0$. Then

$$f'(r_+) = \frac{2d}{r_+^2}. \tag{165}$$

If a nearby continuous numerical profile has a selected root near this simple root, the local inverse estimate gives

$$|\delta r_+| \le \frac{r_+^2}{2d}\,\|\delta f\|_\infty + O(\|\delta f\|_\infty^2). \tag{166}$$

The estimate describes the nearby root and presupposes its localization to a region where the exact derivative is nonzero. It does not make every zero of an arbitrarily oscillatory perturbation physically relevant. The factor is $2.133333333\ldots$ in the benchmark and diverges as $d \downarrow 0$.

At extremality $f(R) = (1 - M/R)^2$ has a double root. A perturbation of size ϵ can move the root by order $\sqrt{\epsilon}$ or remove it entirely. Finite canonical data at a horizon therefore do not imply uniformly well-conditioned horizon detection. The first trapped sphere in the scalar calculation below also occurs at a spatial compactness maximum, and must not be assigned this simple-root estimate without checking its event geometry.

16 Finite radial reconstruction and its stability boundary

On interior points of a uniform radial grid, use the fourth-order derivative

$$(D_4 u)_j = \frac{-u_{j+2} + 8u_{j+1} - 8u_{j-1} + u_{j-2}}{12h}. \tag{167}$$

Its stencil radius is two. Given M, Q, and a fixed branch ε, a radius array satisfying $D_4 R_j > 0$ and $2MR_j - Q^2 > 0$ determines

$$\Pi_h[R]_j = \left(R_j, D_4 R_j, \varepsilon\sqrt{2MR_j - Q^2}, \frac{M D_4 R_j}{\varepsilon\sqrt{2MR_j - Q^2}}, Q \right). \tag{168}$$

This is a local reconstruction of $(R, \Lambda, P_\Lambda, P_R, Q)$, rather than an orthogonal Hilbert-space projector. If it is extended to a full state by discarding the other variables and retaining R, it is idempotent on its admissible image.

Proposition 16.1 (Exact finite reduced constraints)**.** *On the image of* (168)*, the four reduced constraints*

$$\begin{aligned} C_{\Lambda,j} &= \Lambda_j - D_4 R_j, & C_{M,j} &= P_{\Lambda,j}^2 - 2MR_j + Q^2, \\ C_{P,j} &= P_{R,j}P_{\Lambda,j} - M\Lambda_j, & C_{Q,j} &= Q_j - Q \end{aligned} \tag{169}$$

vanish in exact arithmetic. The discrete mass coordinate obtained by replacing R' with $D_4 R$ in (148) *equals M at every reconstructed node.*

Proof. Each constraint vanishes by direct substitution. In the discrete mass coordinate $D_4 R/\Lambda = 1$ and $P_\Lambda^2 = 2MR - Q^2$, leaving $\frac{R}{2}(2M/R - Q^2/R^2) + Q^2/(2R) = M$. □

These identities are exact algebraic constraints on the reconstruction manifold. The original ADM expressions evaluated with finite differences have a different status: the stencil does not obey a pointwise chain rule. Thus $D_4\sqrt{2MR - Q^2}$ need not equal $MD_4R/\sqrt{2MR - Q^2}$ exactly. The original momentum and Hamiltonian constraints are retained as continuum-consistency diagnostics, rather than being relabeled as exact finite invariants.

16.1 A level-independent reconstruction estimate

Assume a uniform regular set

$$R_j \geq R_* > 0, \quad 2MR_j - Q^2 \geq p_*^2 > 0, \quad 0 < d_* \leq D_4 R_j \leq d^*. \tag{170}$$

For two arrays in this set, with the same (M, Q, ε), differentiation along the segment between them gives

$$\|\delta\Lambda\|_\infty \leq \|D_4 \delta R\|_\infty , \tag{171}$$

$$\|\delta P_\Lambda\|_\infty \leq \frac{M}{p_*} \|\delta R\|_\infty , \tag{172}$$

$$\|\delta P_R\|_\infty \leq \frac{M}{p_*} \|D_4 \delta R\|_\infty + \frac{M^2 d^*}{p_*^3} \|\delta R\|_\infty . \tag{173}$$

For example, $\partial_R p = M/p$ and the two partial derivatives of $M\Lambda/p$ have magnitudes at most M/p_* and $M^2 d^*/p_*^3$. The admissible segment preserves the displayed linear radius and derivative bounds. The constants are independent of h.

The input norm contains a radial derivative. A small pointwise radius error alone therefore does not establish a small error in every reconstructed canonical field. This estimate transfers an already controlled discrete $W^{1,\infty}$ radius error to the reconstructed state. It does not create that derivative control from an L^2 estimate.

16.2 Direct and split radius evolution

On the positive branch, restriction of the radius equation to the reconstruction manifold gives

$$\dot{R}_j = g(R_j) + b_j (D_4 R)_j, \qquad g(R) = -\frac{\sqrt{2MR - Q^2}}{R}, \tag{174}$$

where $b = N^r$ is prescribed in the electrovac benchmark. A split alternative is

$$\dot{R} = g(R) + A_h[b]R, \qquad A_h[b]u = \frac{1}{2}(bD_4u + D_4(bu) - (D_4b)u). \tag{175}$$

On smooth data these formulas differ by half of the discrete product-rule defect, which is fourth order away from boundaries. Their exact finite-grid operators differ.

The archived implementation advances the radius explicitly with RK4. The map Π_h can reconstruct the other fields from each admissible stage radius, and its reduced identities hold on that reconstructed image. At output, however, the code overwrites two canonical boundary nodes on each side with exact manufactured data. Those boundary nodes are not generally the exact image of the finite derivative reconstruction. Its reported constraint norms exclude four nodes from each end. This procedure does not use an implicit discrete-gradient slab solve. Conservation by an implicit invariant-preserving method elsewhere in this monograph remains conditional on actually solving its nonlinear step equation. It cannot be used as evidence that this explicit radius algorithm preserves the original ADM Hamiltonian exactly.

Proposition 16.2 (Split-form energy identity)**.** *Let D be skew adjoint for the grid inner product $\langle u, w\rangle_h$, and let multiplication by the real array b be self-adjoint. Then*

$$\langle u, A_h[b]u\rangle_h = -\frac{1}{2}\langle u, (Db)u\rangle_h . \tag{176}$$

For two solutions of the split equation with the same prescribed b, remaining in a positive-radicand interval where $|g'| \le L_g$, one obtains

$$\left\|R(t) - \widetilde{R}(t)\right\|_h \le e^{(L_g + S_1/2)t}\left\|R(0) - \widetilde{R}(0)\right\|_h , \qquad \|Db\|_\infty \le S_1. \tag{177}$$

Proof. Skew adjointness cancels $\langle u, bDu\rangle_h + \langle u, D(bu)\rangle_h$. The remaining term gives (176). With $e = R - \widetilde{R}$, the nonlinear difference satisfies $|g(R_j) - g(\widetilde{R}_j)| \le L_g|e_j|$, where

$$g'(R) = \frac{MR - Q^2}{R^2\sqrt{2MR - Q^2}}. \tag{178}$$

Thus $\frac{1}{2}\partial_t \|e\|_h^2 \le (L_g + S_1/2)\|e\|_h^2$, and integration gives the claim. □

The skew-adjoint hypothesis needs care on a physical radial interval. A periodic D_4 satisfies it, but a globally periodic radius array cannot also satisfy $D_4R_j > 0$ everywhere: summing the derivative over one period gives zero. The periodic identity is an exact algebraic check of the split stencil. It does not by itself give a nonvacuous stability theorem for the globally monotone spherical reconstruction.

For an interval summation-by-parts derivative with a diagonal positive norm matrix H and $HD + D^TH = B$, multiplication by b commutes with H. The uncancelled term is instead

$$\langle u, A_h[b]u\rangle_H = \frac{1}{2}u^TB\ \mathrm{diag}(b)u - \frac{1}{2}\langle u, (Db)u\rangle_H . \tag{179}$$

Its boundary flux must be controlled using an actual boundary closure and appropriate boundary data. Saying that a stencil is centered in the interior does not establish this property for the complete interval operator.

The periodic Fourier symbol is

$$\widehat{D}_4(\theta) = \frac{i}{h}\frac{8\sin\theta - \sin 2\theta}{6}. \tag{180}$$

Its maximum dimensionless magnitude is $1.3722219798\ldots$. Applying the RK4 imaginary-axis stability interval to frozen constant advection gives

$$\frac{|b|\Delta t}{h} \le 2.0612023174\ldots. \tag{181}$$

This is a frozen-coefficient diagnostic. A full variable-coefficient nonlinear RK4 convergence theorem still requires a uniform stability argument for the implemented boundary closure and a suitable stronger norm. Fourth-order consistency, the semidiscrete L^2 estimate, and this one CFL number do not alone prove fourth-order $W^{1,\infty}$ convergence.

17 What the electrovac benchmark measures

The retained run uses

$$F(t,r) = r + 0.08\sin(0.7t)\sin^2\left(\pi\frac{r-0.6}{5.4}\right), \quad 0 \le t \le 0.4, \tag{182}$$

with $M = 1$, $Q = 0.8$, and $\Delta t/h$ close to 0.1. The sampled minimum radial Jacobian is 0.9871380683 and the minimum momentum radicand is 0.56. The exact shift, ghost values, and two boundary nodes on either side are supplied from the manufactured family. Errors are measured on interior nodes after excluding four points from each end.

The recorded state errors are

N	$\|\delta R\|_\infty$	$\|\delta\Lambda\|_\infty$	$\|\delta P_\Lambda\|_\infty$	$\|\delta P_R\|_\infty$
64	$8.3247\,10^{-9}$	$4.1362\,10^{-8}$	$5.7367\,10^{-9}$	$2.3732\,10^{-8}$
128	$5.0495\,10^{-10}$	$2.5079\,10^{-9}$	$3.3079\,10^{-10}$	$1.3795\,10^{-9}$
256	$3.1077\,10^{-11}$	$1.5438\,10^{-10}$	$2.0248\,10^{-11}$	$8.4854\,10^{-11}$
512	$1.9296\,10^{-12}$	$1.0485\,10^{-11}$	$1.2559\,10^{-12}$	$5.3844\,10^{-12}$

The least-squares orders are respectively 3.99856, 3.96010, 4.02393, and 4.00788. The ordinary ADM Hamiltonian on exact projected data has fitted order 3.70410. These are results for a smooth, forced trajectory on four grids. The tiny final errors are close enough to roundoff that they should not be interpreted as arbitrarily extendible asymptotic scaling laws.

On the tested interior of the $N = 512$ reconstructed grid, the retained mass residual is $4.44\,10^{-16}$, the flat-slice residual is zero, the algebraic mass constraint is $2.22\,10^{-15}$, and the reduced momentum residual is $1.11\,10^{-16}$. These floating-point numbers are consistent with the exact algebraic reconstruction theorem. They are not an exact-zero assertion about floating arithmetic or the overwritten boundary nodes.

The negative control evolves all four canonical fields without the reconstruction. At $N = 128$, it reports mass drift $2.25125\,10^{-4}$, radial momentum residual $1.27688\,10^{-3}$, and Hamiltonian residual $1.25735\,10^{-2}$. This gives the reduced reconstruction a meaningful comparator. It does not show that every other constrained algorithm fails or that the direct ADM system is ill posed.

A smooth perturbation of amplitude 10^{-8} evolved with the split form has recorded L^2 amplification between 1.03345462 and 1.03348106 on $N = 128, 256, 512, 1024$. The spread is $2.64438\,10^{-5}$. This supports the bounded experiment's level-insensitive behavior. It leaves the interval boundary and stronger-norm theorem obligations visible.

The outer horizon remains at the fixed areal radius 1.6. Its computational coordinate moves because F changes. The retained final coordinate comparison at $N = 256$ gives $|r_{+,h} - r_{+}^{\text{exact}}| = 1.89262\,10^{-11}$. This measures tracking of a known horizon under coordinate deformation. It does not measure the formation or growth of a black hole.

The electrovac verification record contains nineteen passed checks. The symbolic identity and archived data structure have supporting audit records, but these are distinct from a replay of the electrovac evolution. The numerical table is retained as recorded evidence. Its environment and command provenance are specified in the verification appendix.

18 A dynamical spherical scalar system

To test dynamical concentration of matter, replace electrovac by a real massless scalar. Use $G = c = 1$ and the rescaled field $\varphi = \sqrt{4\pi}\,\phi_{\text{phys}}$. The scalar canonical pair is (φ, P_φ), and the matter additions to the Hamiltonian and momentum constraints are

$$\mathcal{H}_{\text{matt}} = \frac{P_\varphi^2}{2\Lambda R^2} + \frac{R^2\varphi'^2}{2\Lambda}, \qquad \mathcal{D}_{\text{matt}} = P_\varphi\varphi'. \tag{183}$$

The gravitational part is (145) with $Q = 0$.

Impose $R = r$ and $\Lambda = 1$, and set

$$N = \sigma > 0, \quad N^r = \sigma v, \quad P_\Lambda = rv, \quad P_\varphi = r^2\Pi, \quad \Phi = \varphi'. \tag{184}$$

Here Π is the scalar momentum per areal area, not the reconstruction map Π_h. Define

$$m = \frac{r}{2}v^2, \qquad v = \sqrt{\frac{2m}{r}}. \tag{185}$$

The inward branch gives

$$ds^2 = -\sigma^2 dt^2 + (dr + \sigma v dt)^2 + r^2 d\Omega^2. \tag{186}$$

The following derivation applies on smooth patches with $r > 0$, $\sigma > 0$, and $v > 0$. These hypotheses keep every division legitimate.

The momentum constraint removes P_R algebraically:

$$P_R = P_\Lambda' - P_\varphi\Phi. \tag{187}$$

Since $m' = P_\Lambda P_\Lambda'/r - P_\Lambda^2/(2r^2)$, inserting this relation into the Hamiltonian constraint gives

$$m' = r^2\mathcal{E}, \qquad \mathcal{E} = \frac{1}{2}(\Pi^2 + \Phi^2) + v\Pi\Phi. \tag{188}$$

Preservation of $\Lambda = 1$ gives

$$0 = \sigma\left(\frac{v}{r} - \frac{P_R}{r}\right) + (\sigma v)' = \sigma r\Pi\Phi + \sigma' v, \tag{189}$$

hence

$$(\log\sigma)' = -\frac{r\Pi\Phi}{v}. \tag{190}$$

The scalar evolution equations are

$$\dot{\varphi} = \sigma(\Pi + v\Phi), \tag{191}$$
$$\dot{\Phi} = \partial_r[\sigma(\Pi + v\Phi)], \tag{192}$$
$$\dot{\Pi} = r^{-2}\partial_r[r^2\sigma(\Phi + v\Pi)]. \tag{193}$$

If $\Phi = \varphi'$ initially and compatible boundary data are imposed, the first two equations preserve that relation in the continuum.

At $v = 0$, the lapse equation requires a separately specified extension. Smoothness of Π and Φ alone does not establish that $r\Pi\Phi/v$ is removable. A vacuum interval with vanishing scalar data can provide an appropriate extension, but that must be part of the argument. The Gaussian benchmark has very small tails rather than exact compact support, and its numerical treatment of small v is described below. This calculation does not prove a regular-center theorem.

18.1 Exact mass flux and compatibility

Put $A = \Pi + v\Phi$ and $B = \Phi + v\Pi$. The remaining Einstein momentum evolution gives

$$\dot{m} = r^2\mathcal{J}, \qquad \mathcal{J} = \sigma AB. \tag{194}$$

One can verify its origin directly from the P_Λ equation. In the flat-slice gauge, the gravitational terms give $\sigma v P_\Lambda' - \sigma v^2/2 - r\sigma'$. The lapse-gradient term is essential. The matter variation contributes $\sigma r^2(\Pi^2 + \Phi^2)/2$. As $\dot{m} = v\dot{P}_\Lambda$, their sum becomes

$$\dot{m} = \sigma v\left[m' + \frac{r^2}{2}(\Pi^2 + \Phi^2)\right] - rv\sigma'. \tag{195}$$

Using (188) and (190) yields $\dot{m} = r^2\sigma[v(\Pi^2 + \Phi^2) + (1 + v^2)\Pi\Phi]$, which is (194).

Proposition 18.1 (Continuum mass-constraint propagation)**.** *On the regular patch, the scalar evolution, lapse relation, and mass-flux evolution satisfy*

$$\partial_t(r^2\mathcal{E}) = \partial_r(r^2\mathcal{J}). \tag{196}$$

Thus $\partial_t(m' - r^2\mathcal{E}) = 0$. A radial mass constraint imposed initially propagates under this continuum system.

Proof. Expanding (193) gives $\dot{\Phi} = (\sigma A)'$ and $\dot{\Pi} = (\sigma B)' + 2\sigma B/r$. Differentiate $\mathcal{E}$ and subtract the radial derivative of the flux. Terms involving Π', Φ', and v' cancel, leaving

$$\partial_t(r^2\mathcal{E}) - \partial_r(r^2\mathcal{J}) = r^2\left[\sigma' AB + \Pi\Phi\dot{v}\right]. \tag{197}$$

Since $m = rv^2/2$, the flux equation implies $\dot{v} = r\sigma AB/v$. The remaining expression is

$$r^2 AB\left[\sigma' + \frac{\sigma r \Pi\Phi}{v}\right] = 0 \tag{198}$$

by (190). This establishes the identity without assuming the radial constraint in its last cancellation step. □

On a fixed radial interval $[r_i, r_o]$, the exact balance is

$$\frac{d}{dt}(m(t, r_o) - m(t, r_i)) = r_o^2\mathcal{J}(t, r_o) - r_i^2\mathcal{J}(t, r_i). \tag{199}$$

Boundary fluxes, rather than a statement that a finite-domain mass must be exactly constant, are the appropriate continuum reference. The numerical prefix solve below does not satisfy this identity algebraically at every finite grid and time step.

18.2 Characteristic channels and the loss of a coordinate energy bound

Define $w_+ = \Pi + \Phi$ and $w_- = \Pi - \Phi$. Their equations are

$$\dot{w}_+ = \sigma(v+1)w_+' + [\sigma v' + \sigma'(v+1)]w_+ + \frac{\sigma}{r}[(v+1)w_+ + (v-1)w_-], \tag{200}$$

$$\dot{w}_- = \sigma(v-1)w_-' + [\sigma v' + \sigma'(v-1)]w_- + \frac{\sigma}{r}[(v-1)w_- + (v+1)w_+]. \tag{201}$$

Written as transport equations $\dot{w} + c\,w' =$ source, their coordinate speeds are

$$c_+ = -\sigma(v+1), \qquad c_- = \sigma(1-v). \tag{202}$$

The plus channel is inward on the positive branch. The minus channel is outward when $v < 1$, stationary at $v = 1$, and inward when $v > 1$.

The same channels give the exact algebraic identities

$$\mathcal{E} = \frac{1}{4}[(1+v)w_+^2 + (1-v)w_-^2], \tag{203}$$

$$\mathcal{J} = \frac{\sigma}{4}[(1+v)^2 w_+^2 - (1-v)^2 w_-^2]. \tag{204}$$

For $0 \le v < 1$, $\mathcal{E} \ge (1-v)(w_+^2 + w_-^2)/4$. This lower bound degenerates at a horizon. For $v > 1$, the outward-channel coefficient in this coordinate mass density is negative. The scalar's local normal-frame quadratic density $(\Pi^2 + \Phi^2)/2$ remains nonnegative. The loss of coercivity in $\mathcal{E}$ therefore diagnoses this slicing's mass-coordinate energy estimate, rather than negative physical scalar energy.

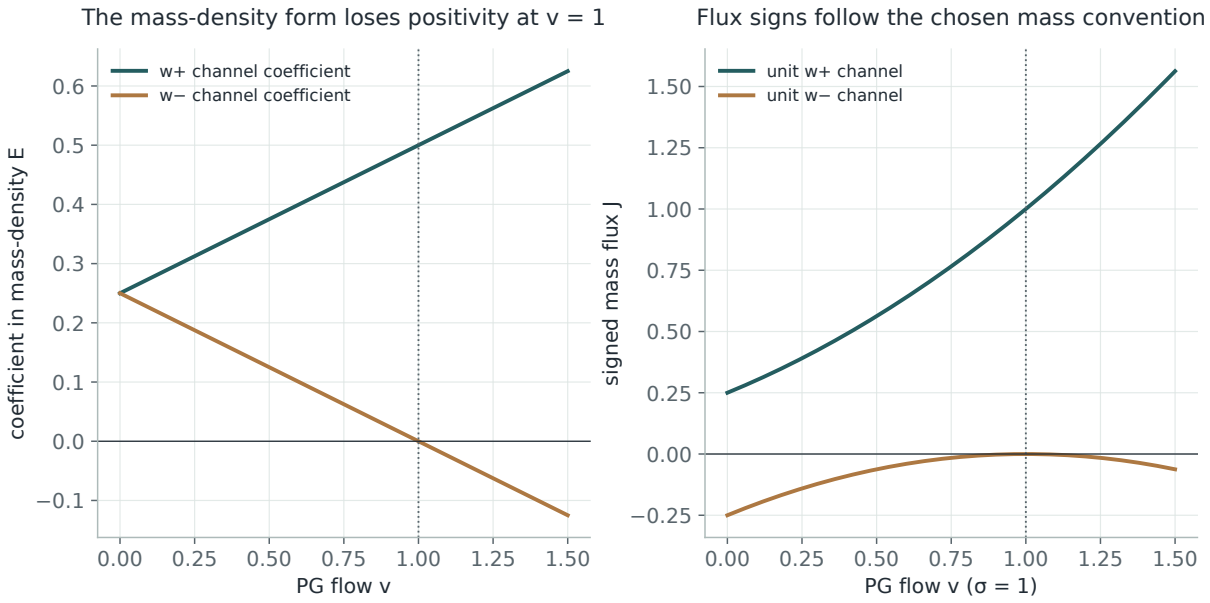


Figure 7. The mass density and mass flux resolve into the two scalar characteristic channels. The outgoing characteristic and its contribution to the coordinate mass density degenerate at $v = 1$. These are exact algebraic profiles, not a newly evolved collapse solution.

19 Trapped spheres and the finite-stencil scalar experiment

For (186), choose future radial null vectors

$$\ell_{\pm} = \sigma^{-1}\partial_t + (-v \pm 1)\partial_r. \tag{205}$$

Because the sphere area is $4\pi r^2$, its expansion along ℓ is $2\ell(r)/r$. Consequently

$$\theta_{+} = \frac{2}{r}(1 - v), \qquad \theta_{-} = -\frac{2}{r}(1 + v). \tag{206}$$

On the positive branch, $v = 1$ is a marginally outer trapped sphere, and $v > 1$ gives both expansions negative. Equivalently, the compactness $2m/r$ equals or exceeds one. The normalization of the null vectors affects expansion magnitude but not these signs. A marginally trapped sphere is a local geometric diagnostic. It is not by itself identification of a global event horizon.

19.1 What is local in the algorithm

At each RK4 stage the matter channels determine the geometry through two radial solves. The mass constraint is integrated outward with $m(r_{\text{exc}}) = 0$. The lapse derivative is integrated and the result normalized by $\sigma(r_{\max}) = 1$. These are triangular prefix operations. A matter value near the inner end can affect the reconstructed mass and lapse far along the interval in the same stage. The full stage map is therefore globally coupled.

The characteristic evolution uses the second-order one-sided derivatives

$$(D_{+}u)_j = \frac{-3u_j + 4u_{j+1} - u_{j+2}}{2h}, \qquad (D_{-}u)_j = \frac{3u_j - 4u_{j-1} + u_{j-2}}{2h}. \tag{207}$$

The inward channel uses D_+ and the outgoing channel uses D_- outside a horizon, switching to D_+ inside. These interior matter stencils have radius two. The source also uses a radius-two fourth-order interior derivative for v' and separate one-sided boundary formulas. Finite stencil radius for these derivatives must not be promoted to a finite dependency cone for the geometry-coupled step.

The mass integration uses radial RK4 with interpolated midpoint fields, followed by Simpson-type lapse quadrature. These are approximate quadratures. They are neither exact mass-constraint algebraic reconstruction nor an implicit invariant-preserving solver. The scalar spatial evolution is second order despite its fourth-order time integrator.

Several implementation conventions are mathematically material. The mass right side and the reconstructed v use $\sqrt{\max(2m/r, 0)}$. The lapse derivative is set to zero where $v \leq 10^{-13}$. Incoming characteristic values are explicitly set after a completed RK4 step, with separate boundary formulas during stages. These choices guard a finite calculation near vacuum and boundaries. An exact continuum proof on $v > 0$ does not establish their uniform consistency as the threshold is varied or $h \downarrow 0$.

19.2 Initial data and boundary scope

The numerical domain is $[r_{\text{exc}}, r_{\text{max}}] = [0.05, 30]$. It replaces a regular center by a small excised core. The characteristic data are $w_-(r_{\text{exc}}) = 0$ and $w_+(r_{\text{max}}) = 0$. These prohibit prescribed incoming radiation from the inner and outer boundaries in the respective channels. The radial mass boundary condition additionally fixes the excised mass to zero at every geometry solve.

The initial scalar packet is

$$\varphi(0, r) = Ae^{-(r-10)^2}, \qquad \Phi(0, r) = \partial_r \varphi(0, r), \qquad \Pi(0, r) = \Phi(0, r) + \frac{\varphi(0, r)}{r}. \tag{208}$$

The last equation is the flat-space ingoing condition for $\chi = r\varphi$, used to initialize the scalar. The nonlinear gravitational constraints are then solved. It is therefore an ingoing packet prescription, rather than exactly freezing the geometry to Minkowski space.

The strong amplitude is $A = 0.10$ and the control amplitude is $A = 0.05$. The strong event occurs near $r = 3.1$, well separated from both boundaries in coordinate distance. That separation is useful but is not a substitute for studying boundary influence: the geometry solve is global, and finite-domain mass drift can contain both boundary and discretization effects.

19.3 Recorded first marginally trapped sphere

The event finder takes a quadratic interpolation of the spatial maximum of $2m/r$ and a linear interpolation between consecutive time steps bracketing its first crossing of one. The archived event table is

N_r	h	t_H	r_H	$\Delta M/M_0$
600	0.0500000	4.7133579	3.0518959	$-1.1205\,10^{-2}$
800	0.0374844	4.7162491	3.0506543	$-1.6427\,10^{-2}$
1000	0.0299800	4.7059219	3.0766099	$-1.0696\,10^{-2}$
1200	0.0249791	4.7004729	3.0905271	$-7.6261\,10^{-3}$
1600	0.0187305	4.6952990	3.1040860	$-4.6604\,10^{-3}$
2000	0.0149825	4.6930194	3.1101678	$-3.3322\,10^{-3}$
2400	0.0124844	4.6918209	3.1134115	$-2.6241\,10^{-3}$

The coarsest two event coordinates are not monotone with refinement. On the six finest resolutions the retained nonlinear fits are

$$t_H(h) \simeq 4.6893522370 + 33.88577338\, h^{2.173853282}, \tag{209}$$

$$r_H(h) \simeq 3.1204965744 - 65.59310603\, h^{2.084394965}. \tag{210}$$

These fit exponents are consistent with second-order event convergence on this bounded family. They are not a convergence proof, a reliable error bar for every observable, or a critical-collapse exponent.

For $N_r = 2400$, the interpolated event has

$$t_H = 4.69182094580, \quad r_H = 3.11341154797, \quad m_H = r_H/2 = 1.55670577399. \tag{211}$$

The relation $m_H = r_H/2$ is imposed by the interpolated marginal-sphere criterion. It is not an independent numerical conservation test.

The stored event fields come from the next completed step, at $t = 4.69206969694$, rather than exactly the interpolated t_H. Their recorded outer mass is 1.72538868709, compared with initial outer mass 1.72992821318, a fractional change of $-2.6241124\,10^{-3}$. The same stored fields have minimum lapse 0.999829382946, lapse at the sampled peak 1.02582462096, and maximum absolute scalar Ricci curvature 7.35797274672. This timing distinction is small but matters when those quantities are compared with an independently interpolated event.

The scalar curvature diagnostic follows the rescaled convention

$$\mathcal{R} = 2(\Phi^2 - \Pi^2). \tag{212}$$

Finite lapse and curvature values show that this calculation can reach the marginal-sphere event without a coordinate blow-up. They do not provide continuation to a singularity or control of subsequent trapped-region evolution.

19.4 Control and independent diagnostics

The weak pulse on $N_r = 1200$ remains untrapped through $t = 5$, with maximum compactness 0.146722872367 and fractional outer-mass change $-4.5106042\,10^{-3}$. It is a finite-time control in an excised domain. It does not establish global dispersion, passage through a regular center, or a critical amplitude separating two complete physical outcomes.

At the smooth pre-horizon time $t = 3$, the $N_r = 1200$ retained relative residuals are

Diagnostic	Relative residual
Radial mass constraint	$2.27319745\,10^{-6}$
Lapse constraint	$2.80774843\,10^{-6}$
Local mass flux against a time difference	$1.26971871\,10^{-3}$

The flux diagnostic is independent of simply inserting the geometry solve's own right side into its output. Its finite time difference is still an approximation. The separately reported absolute outer-mass flux-balance defect is 0.00230311127. A small total outer-mass change cannot be attributed solely to physical boundary flux unless that flux is integrated and its numerical balance is controlled.

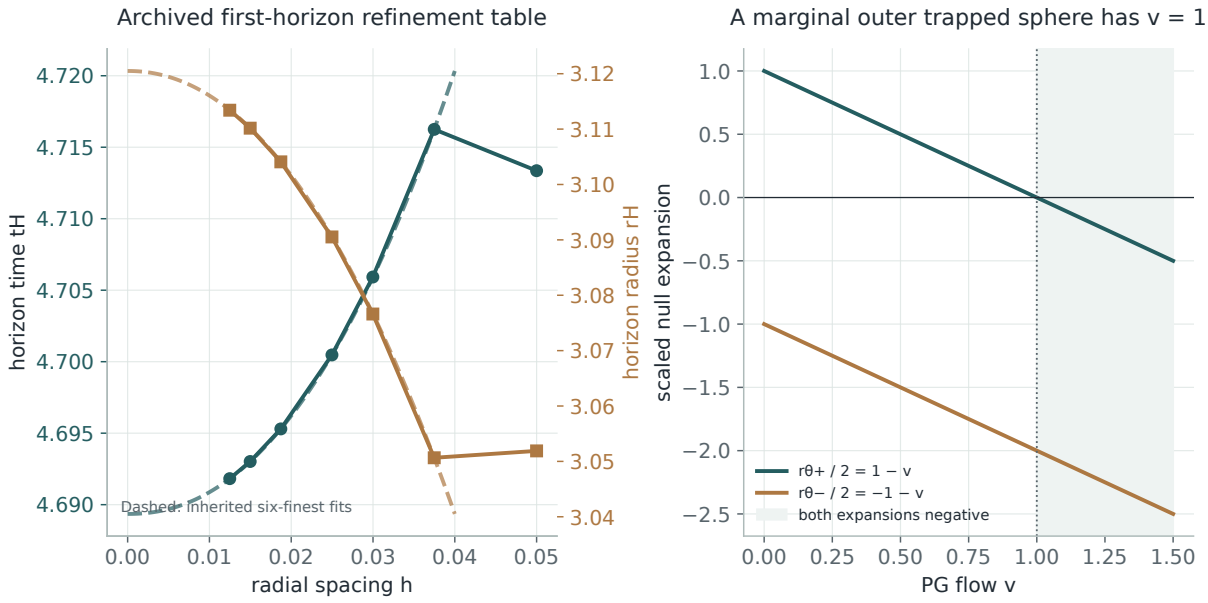


Figure 8. Archived event coordinates under radial refinement, together with the exact expansion and characteristic sign change at $v = 1$. The event table is replotted from retained data. No new scalar-collapse evolution was performed for this figure.

The scalar verification record contains twelve passed checks. Two concern symbolic identities. The remaining checks are bounded numerical tests with declared thresholds. The symbolic audit, null-channel algebra, and seven-row data structure have separate supporting records. Those checks do not constitute a replay of the full collapse or an authentication of its original execution environment. The verification appendix states that evidence boundary.

20 What remains beyond these strong-field sectors

There are exact results here: the electrovac mass identity, its regular-patch constraint equivalence, the reconstructed finite reduced constraints, the scalar flat-slice reduction, the continuum local mass flux, and its compatibility with the radial mass constraint. Their proofs use the displayed chart and symmetry hypotheses. Neither numerical rounding nor a missing fresh Lean build changes them into a machine-checked unrestricted field theory.

There are also controlled numerical observations: a prescribed electrovac gauge trajectory crossing its known outer horizon, and an excised spherical scalar experiment reaching a first marginally trapped sphere. Their retained refinement data and negative control make useful, bounded tests. The next proofs must address the actual boundary closure, regularity margins, vacuum extension, solver conventions, and errors in the full coupled stage map.

For electrovac, a nonvacuous interval energy theorem must control boundary flux and then establish the stronger norm needed by the reconstruction estimate. A full RK4 theorem needs a variable-coefficient argument, rather than the frozen advection symbol alone. For scalar matter, the mass and lapse prefix solves need stability and consistency estimates coupled to the characteristic evolution, including behavior near $v = 0$ and the excision boundary. Event convergence needs quantitative control

of the compactness maximum and its temporal crossing. A tangential spatial root cannot inherit a simple-root horizon bound automatically.

Beyond these radial sectors, finite-range nonlinear gravitational closure remains a separate problem. Spherical symmetry removes the radiative gravitational degrees of freedom. A general $3 + 1$ Einstein–Maxwell–Dirac construction must restore them, choose boundary and gauge conditions, establish the relevant constraint algebra and chart regularity, and prove consistency and stability of its actual evolution. Numerical marginal-sphere formation here establishes no global cosmic-censorship result, singularity theorem, nonspherical stability theorem, or unique event-horizon reconstruction.

The exact quasienergy chapter in the source corpus supplies a further boundary of interpretation. Its componentwise momentum change of variables conjugates a branchwise nonlinear Lorentz action to ordinary Lorentz transformations. The resulting momentum-space metric is a pullback of a flat metric. This is a kinematic coordinate construction, not a solution of the gravitational constraints. Likewise, positive curvature of a low-momentum dispersion relation is not spacetime curvature. A gravity coupling of the quasienergy system requires its own tetrad, connection, stress-energy, and dynamical constraint analysis. That obligation is not filled by relabeling a momentum reparametrization as strong-field geometry.

21 Constraint structure and the limits of reconstruction

The gravitational branch now has a continuous account from transport geometry to constrained dynamics and spherical tests. The geometric examples identify the additional assumptions needed when a propagation rule is interpreted as gravity. Reduced Gowdy evolution and finite linearized ADM theory make constraint propagation concrete. The nonlinear product-rule obstruction then identifies a specific limitation of a finite pointwise discretization.

The alternatives address that limitation in their own settings. Matrix moment maps retain exact finite algebraic closure. Homological transfer requires a genuine fine differential and controlled contraction data. Resolved quadrature retains the stated polynomial identities under its degree budget, while nonlinear evolution can require additional spectral bands. These constructions therefore carry distinct state spaces and comparison problems.

The Lorentzian classical action supplies a gauge-theoretic parent. Its boundary description and local second-class reduction require the declared regularity, nondegeneracy, canonical transformation, and invertibility hypotheses. Projected discrete-gradient evolution preserves the stated autonomous invariants for an exact step solve within that chart. None of these hypotheses can be discarded when changing coordinates or moving toward a singular configuration.

Spherical mass and flux identities provide unusually direct checks. Electrovac benchmarks test a prescribed slicing trajectory of a fixed charged geometry. The scalar-collapse implementation adds evolved matter and radial constraint solves. Its trapped-sphere diagnostic is a bounded numerical result for the declared data and implementation. Establishing an event in a controlled continuum solution requires an additional existence and convergence argument.

21.1 Mathematical work that remains

A useful next nonlinear theorem should identify a compatible sequence of state spaces and maps, give a level-uniform stability estimate, and state the topology in which the constraint and evolution defects tend to zero. Exact evaluation at one resolution does not itself provide that sequence. A reduction

theorem that crosses charts must also control degeneration of the gauge-fixing and Dirac matrices, boundary conditions, and the change of canonical variables.

The strong field route needs estimates adapted to its actual constraint solve and boundary closure. For scalar matter, this includes the nonlocal radial dependence of reconstructed mass and lapse, the branch of the slicing variables, and convergence strong enough to control the trapped-sphere diagnostic. Extending the construction to nonspherical geometry introduces gravitational degrees of freedom absent from a spherical model.

The inherited massive fermion conventions also need a physical comparison in flat space. Lorentzian signature, the Clifford relation, the kinetic factor, and the adjoint must be matched together. The algebraic gauge model does not settle that normalization.

A quantum theory requires further choices and arguments. The classical master equation does not define a quantum measure or settle anomalies. A common quantum geometry–matter construction and a controlled semiclassical comparison remain research obligations. The verification companion lists these obligations beside the source results they extend.

21.2 From geometric to arithmetic refinement

The next chapter develops register factorization, Haar packets, spectral reduction, and infinite-depth quantum hierarchies. It changes the model while preserving the discipline used here: finite identities, refinement compatibility, and analytic limits require separate arguments. The covariant fiber and cosmology chapters then examine whether a chosen refinement sector admits a physically controlled production history and consistent observable transport.

> **A concrete next comparison**
>
> Name the phase space, constraints, evolution, boundary conditions, and map between resolutions. State the regularity and chart on which the map is defined. Then prove the stability and convergence needed for the physical conclusion being sought.

Sources and further reading

Project sources

[P1] Komorebi P. *Projector Refinement Physics, Volume I*. The source studies cover weak fields, holonomy, Gowdy reduction, linearized ADM constraints, nonlinear transfer, resolved connections, the Lorentzian parent, regular reduction, electrovac geometry, and scalar collapse.

[P2] Komorebi P. *Projector Refinement Physics, Volume I: Technical Supplement*. The nonlinear and strong-field claim dossiers specify boundary, spectral, and numerical hypotheses. The verification appendix maps these hypotheses to the constructions used here.

[P3] Komorebi P. *Komorebi Physics formal source*. Its projector, gauge, and free Dirac results provide related mathematical infrastructure. They do not constitute a compiler certificate for the nonlinear gravity and collapse statements.

Primary context

[C1] R. Arnowitt, S. Deser, and C. W. Misner. "The Dynamics of General Relativity." In L. Witten, ed., *Gravitation: An Introduction to Current Research*, Wiley, 1962, Chapter 7, pp. 227–265. Author republication arXiv:gr-qc/0405109v1, 19 May 2004. https://arxiv.org/abs/gr-qc/0405109v1.

[C2] A. Ashtekar. "New Variables for Classical and Quantum Gravity." *Physical Review Letters* **57** (1986), 2244–2247. https://doi.org/10.1103/PhysRevLett.57.2244.

[C3] I. A. Batalin and G. A. Vilkovisky. "Gauge algebra and quantization." *Physics Letters B* **102** (1981), 27–31. https://doi.org/10.1016/0370-2693(81)90205-7.

[C4] A. S. Cattaneo, P. Mnev, and N. Reshetikhin. "Classical BV theories on manifolds with boundary." *Communications in Mathematical Physics* **332** (2014), 535–603. Author preprint arXiv:1201.0290v3, 28 June 2012. https://arxiv.org/abs/1201.0290v3; https://doi.org/10.1007/s00220-014-2145-3.

[C5] A. S. Cattaneo, P. Mnev, and N. Reshetikhin. "Perturbative quantum gauge theories on manifolds with boundary." *Communications in Mathematical Physics* **357** (2018), 631–730. Author preprint arXiv:1507.01221v2, 30 May 2016. https://arxiv.org/abs/1507.01221v2; https://doi.org/10.1007/s00220-017-3031-6.

[C6] R. I. McLachlan, G. R. W. Quispel, and N. Robidoux. "Geometric integration using discrete gradients." *Philosophical Transactions of the Royal Society A* **357** (1999), 1021–1045. https://doi.org/10.1098/rsta.1999.0363.

[C7] C. W. Misner and D. H. Sharp. "Relativistic Equations for Adiabatic, Spherically Symmetric Gravitational Collapse." *Physical Review* **136** (1964), B571–B576. https://doi.org/10.1103/PhysRev.136.B571.

[C8] J. Ziprick and G. Kunstatter. "Numerical study of black-hole formation in Painlevé–Gullstrand coordinates." *Physical Review D* **79** (2009), 101503(R). Author preprint *Spherically Symmetric Black Hole Formation in Painlevé–Gullstrand Coordinates*, arXiv:0812.0993v1, 4 December 2008. https://arxiv.org/abs/0812.0993v1; https://doi.org/10.1103/PhysRevD.79.101503.

How these sources are used

The ADM republication supplies the canonical setting for lapse, shift, and constraints. Ashtekar and Batalin–Vilkovisky supply the historical connection and master-equation settings. These references do not establish the projector model's continuum convergence or its quantum anomaly cancellation.

The two Cattaneo–Mnev–Reshetikhin preprints provide the precise comparison for boundary reduction and BV elimination. Section 3.1.3 of arXiv:1201.0290v3 constructs boundary data through a presymplectic quotient and its projectability conditions. Bulk gauge invariance alone therefore does not discharge the missing regular boundary construction in the present gravity model. In arXiv:1507.01221v2, Theorem 2.9 and Corollary 2.12 give BV pushforward and preservation of the quantum master equation for compatible finite-dimensional odd-symplectic data. The identity is conditional on the fine quantum master equation and well-defined integration. It does not manufacture an anomaly-free quantum gravity measure.

McLachlan–Quispel–Robidoux is the established discrete-gradient context. The preservation identity used here is proved directly from the segment integral and skewness of the endpoint Dirac tensor. Misner–Sharp identifies the spherical mass tradition. The charged mass-coordinate identity in this volume is independently differentiated in the stated canonical convention. Ziprick–Kunstatter is prior work on scalar collapse in Painlevé–Gullstrand coordinates. Its regular-center and critical-collapse analyses are distinct from the excised, first-event benchmark retained here.

Relevant full-text sections of the ADM, both Cattaneo–Mnev–Reshetikhin, and Ziprick–Kunstatter author preprints were inspected in their stated versions. For Ashtekar, Batalin–Vilkovisky, McLachlan–Quispel–Robidoux, and Misner–Sharp, bibliographic metadata was authenticated, but publisher full-text access was unavailable. Their uninspected results are not used to discharge a proof obligation. The source record describes this focused comparison and its limits. No priority claim or exhaustive novelty review is made.

PART III

Arithmetic and polyphase refinement

Arithmetic registers provide another exact resolution of a quantum state. Compatibility determines when refinement preserves the dynamics. Eliminating hidden channels produces self-energy and memory, while infinite depth requires its own analytic domain.

IN THIS PART

Adic and Polyphase Quantum Refinement 161

Adic and Polyphase Quantum Refinement

Normalized packets, spectral elimination, and compatible limits

Arithmetic resolution and quantum refinement

The preceding constructions compare geometric and spectral resolutions. Arithmetic refinement provides another concrete family: finite unit registers, Haar packets, hierarchical Hamiltonians, and polyphase channels. The question remains the same. Which maps preserve the state, which preserve the dynamics, and what additional structure appears when hidden components are eliminated?

Normalized Haar averages identify coarse states inside a fine Hilbert space. Compatibility of the Hamiltonian, rather than normalization alone, determines whether that subspace carries the same evolution. Schur reduction and its time-domain memory equation describe the effect of components that cannot simply be discarded.

Infinite depth introduces a separate analytic problem. The hierarchical model specifies a common Hilbert space, a maximal operator domain, and the topology of approximation. Spectral decimation and hypergeometric companion functions provide exact coordinates within their declared domains. None of these identities alone identifies arithmetic depth with a physical observable.

The following account of covariant fiber dynamics supplies a candidate physical coupling and examines its energy and production constraints. That transition separates a mathematical refinement structure from a proposed model of matter.

The visual and verification appendices distinguish exact identities, finite computations, and formal coverage. The infinite arithmetic hierarchy is presented with its written proofs and is not assigned a Lean certificate beyond the statements specifically identified there.

Reading routes and conventions

The finite route begins with unit coordinates and Haar normalization. It continues through exact intertwining and controlled leakage. The spectral route begins with the block resolvent and follows radial chains, residue channels, and Chebyshev elimination. The limiting route begins with cylinder projections, constructs the maximal hierarchical Hamiltonian, and then treats compatible tensor dynamics.

An embedding and an elimination have different jobs. An isometry retains the old inner product. A unitary polyphase transformation retains the full fine operator in a new basis. A Schur complement solves for hidden coordinates on a regular spectral chart. A projection approximation converges only in the topology established by its proof. These operations should be identified before their consequences are compared.

Convention	Meaning
Hilbert normalization	Finite registers use an explicitly stated orthonormal basis. Cylinder functions use probability Haar measure. The conversion between counting coordinates and normalized indicators includes the square root of the cylinder measure.
Refinement	An isometry V has $V^*V = I$. Its old-range projector is $P = VV^*$ and its new-detail projector is $Q = I - P$. Exact dynamical compatibility is a separate identity, $H'V = VH$.
Hierarchy	The finite chapter introduces cumulative projectors E_r and their details. The infinite chapter writes E_m for conditional expectations and $\mathsf{F}_m = \mathsf{E}_m - \mathsf{E}_{m-1}$ for depth details, with the constant sector declared separately.
Self-adjointness	A finite Hamiltonian is Hermitian. Complex coupling amplitudes appear with conjugates in the reverse direction. An unbounded Hamiltonian also requires a dense domain and a proof that the declared realization is self-adjoint.
Spectral parameter	Schur formulas use $z - H$ and require the hidden resolvent $(z - QHQ)^{-1}$. Real-energy formulas are restricted to their regular chart. They are not evaluated at hidden poles without an extension argument.
Spectral coordinates	A residue or polyphase transformation is a change of basis. Selecting a channel, tracing a subsystem, averaging amplitudes, and eliminating hidden coordinates are different operations.
Limits	Strong vector, strong operator, graph-norm, and operator-norm convergence have distinct quantifiers. The infinite chapter proves the stated limits for its compatible block model.
Evidence	Ordinary derivations, fresh bounded calculations, retained program outputs, and archived formal source are recorded separately. A small residual does not prove a universal infinite-depth or physical claim.

The technical source concordance identifies statements by exact source member, formula, and claim. Each result here carries its own hypotheses. A source title or a recorded test result is not a substitute for those hypotheses.

Finite registers and normalized refinement

1 Finite unit registers and protected coordinates

Adic refinement first changes the address of a state. Its amplitudes remain complex numbers, and its measurement probabilities remain ordinary Born probabilities. For a finite address set X we use the counting inner product on $\ell^2(X;\mathbb{C})$, so the address vectors $|x\rangle$ are orthonormal. A bijection of address sets therefore gives a permutation unitary without changing the quantum probability rule. Arithmetic determines the address factorization. A Hamiltonian and a measurement prescription supply the additional physical choices.

The first coordinate issue is specific to units. For an odd prime p and an integer $h \geq 0$, set

$$U_{p,h} = (\mathbb{Z}/p^{h+1}\mathbb{Z})^{\times}.$$

For $p = 2$ the useful indexing starts one digit later:

$$U_{2,h} = (\mathbb{Z}/2^{h+2}\mathbb{Z})^{\times}.$$

The shifted dyadic index retains the sign even when the principal phase register has dimension one. It prevents the two units modulo four from being collapsed into one phase coordinate.

1.1 Torsion and principal units

Proposition 1.1 (Finite unit coordinates)**.** *For odd p, choose a unit $\zeta_{p,h}$ of order $p-1$ whose reduction generates $(\mathbb{Z}/p\mathbb{Z})^{\times}$. The maps*

$$(t,a) \longmapsto \zeta_{p,h}^{t}(1+p)^{a} \pmod{p^{h+1}}, \qquad (t,a) \in \mathbb{Z}/(p-1)\mathbb{Z} \times \mathbb{Z}/p^{h}\mathbb{Z}, \tag{1}$$

$$(s,a) \longmapsto (-1)^{s}5^{a} \pmod{2^{h+2}}, \qquad (s,a) \in \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2^{h}\mathbb{Z} \tag{2}$$

are group isomorphisms onto $U_{p,h}$ and $U_{2,h}$ respectively.

Proof. For odd p the valuation identity

$$v_p((1+p)^{p^k} - 1) = k + 1 \qquad (k \geq 0)$$

shows that $1 + p$ has order p^h modulo p^{h+1}. To check the identity, the case $k = 0$ is immediate. If $u = 1 + cp^{k+1}$ with $p \nmid c$, the binomial expansion of $u^p - 1$ has first term cp^{k+2}. Every remaining term has strictly larger p-valuation when p is odd. This proves the induction. The subgroup generated by $1 + p$ is contained in the kernel of reduction modulo p and has the same cardinality p^h, so it is that entire kernel.

For completeness, an order-$p-1$ lift of any nonzero residue can be chosen coherently. The root equation $x^{p-1} = 1$ has nonzero derivative $(p-1)x^{p-2}$ modulo p. If x_n is a root modulo p^n, substituting $x_n + cp^n$ modulo p^{n+1} determines one and only one residue c that cancels the error. Starting with a generator of $(\mathbb{Z}/p\mathbb{Z})^{\times}$ produces $\zeta_{p,h}$ with the required order. The cyclicity of this residue group follows from the elementary fact that a degree-d polynomial over a field has at most d roots: in a finite subgroup of the multiplicative group, the least common multiple m of the element orders is realized by an element. To see this, take for each prime divisor of m a component having the maximal prime-power order and multiply those commuting components of coprime orders. All elements solve $x^m = 1$.

Consequently the subgroup has size at most m, and Lagrange's theorem gives the reverse bound. The subgroup is cyclic.

If $\zeta_{p,h}^{t}(1+p)^{a}=1$, reduction modulo p gives $t=0$ in $\mathbb{Z}/(p-1)\mathbb{Z}$. The principal-unit order then gives $a=0$ in $\mathbb{Z}/p^{h}\mathbb{Z}$. Thus the displayed map is injective, and its domain has size $(p-1)p^{h}=|U_{p,h}|$, which proves surjectivity.

For the dyadic case, successive squaring gives

$$v_2(5^{2^k}-1)=k+2.$$

Indeed $5-1$ has valuation two, and for every such power the factor $5^{2^k}+1$ has valuation one. Thus 5 has order 2^h modulo 2^{h+2}. Its powers are exactly the 2^h units congruent to one modulo four. Multiplication by -1 supplies the units congruent to three modulo four. These two classes are disjoint, giving the claimed unique coordinates. □

In the odd-prime register t is a torsion coordinate and a is a principal phase coordinate. In the dyadic register s is a sign bit and a is again the principal phase. The Hilbert-space factorizations are therefore

$$\begin{aligned}\ell^2(U_{p,h}) &\cong \mathbb{C}^{p-1}_{\text{tors}}\otimes\ell^2(\mathbb{Z}/p^h\mathbb{Z}), && p \text{ odd},\\ \ell^2(U_{2,h}) &\cong \mathbb{C}^{2}_{\text{sign}}\otimes\ell^2(\mathbb{Z}/2^h\mathbb{Z}).\end{aligned} \tag{3}$$

These are unitary changes of basis. They do not assert that a generic state is a product state, or that a generic Hamiltonian is a sum of independent register Hamiltonians. Entanglement and interactions are preserved under the change of coordinates.

The lifting construction in the proof permits the coherent choice $\zeta_{p,h+1}\equiv\zeta_{p,h}\pmod{p^{h+1}}$. With that choice, reduction of unit addresses retains t and reduces a modulo p^h. If the torsion lifts are instead chosen independently at each depth, the finite factorizations remain exact, but a comparison between levels must include the resulting coordinate changes. A levelwise permutation is not by itself a specified refinement tower.

1.2 CRT readdressing and the Fourier dual

Let $\mathcal{P}_{\text{o}}$ be a finite set of odd primes and put

$$M=2^{h_2+2}\prod_{p\in\mathcal{P}_{\text{o}}}p^{h_p+1}.$$

The Chinese remainder theorem gives a bijection from the unit addresses modulo M to their local unit addresses. Combining it with Proposition 1.1 gives a permutation unitary

$$\mathcal{U}_{\text{CRT}}:\ell^2\big((\mathbb{Z}/M\mathbb{Z})^{\times}\big)\longrightarrow\mathbb{C}^{2}_{\text{sign}}\otimes\ell^2(\mathbb{Z}/2^{h_2}\mathbb{Z})\otimes\bigotimes_{p\in\mathcal{P}_{\text{o}}}\left(\mathbb{C}^{p-1}_{\text{tors}}\otimes\ell^2(\mathbb{Z}/p^{h_p}\mathbb{Z})\right). \tag{4}$$

For each basis vector, $\mathcal{U}_{\text{CRT}}|u\rangle$ is the basis vector with its unique tuple of coordinates. Distinct units have distinct tuples, so inner products are preserved. The dimensions also agree:

$$\varphi(M)=2^{h_2+1}\prod_{p\in\mathcal{P}_{\text{o}}}(p-1)p^{h_p}.$$

Multiplication of units becomes addition in each cyclic coordinate.

We use the positive-phase Fourier convention

$$(F_n)_{kx} = n^{-1/2} \exp(2\pi \mathrm{i} kx/n). \tag{5}$$

The character transform of the unit group consequently factors as

$$F_{U_M} = \mathcal{U}_{\mathrm{CRT}}^{*} \left[F_2 \otimes F_{2^{h_2}} \otimes \bigotimes_{p \in \mathcal{P}_{\mathrm{o}}} (F_{p-1} \otimes F_{p^{h_p}}) \right] \mathcal{U}_{\mathrm{CRT}}. \tag{6}$$

Here the unit-group character labels have been identified with the same cyclic coordinate sets. This is a transform on the multiplicative unit group, rather than the additive Fourier transform on all residues modulo M.

The additive transform has a related factorization, with different permutations on its input and its dual output. Let a, b be positive coprime integers and $N = ab$. Define

$$\mathcal{P}_{\mathrm{in}} |n\rangle = |n \bmod a\rangle \otimes |n \bmod b\rangle,$$
$$\mathcal{P}_{\mathrm{out}} |k\rangle = |kb^{-1} \bmod a\rangle \otimes |ka^{-1} \bmod b\rangle,$$

where the inverses are taken in the indicated residue rings. Then

$$F_N = \mathcal{P}_{\mathrm{out}}^{*} (F_a \otimes F_b) \mathcal{P}_{\mathrm{in}}. \tag{7}$$

To verify the phase, let $n_a = n \bmod a$ and $n_b = n \bmod b$. At the dual coordinates $k_a = kb^{-1} \bmod a$ and $k_b = ka^{-1} \bmod b$, the product kernel is

$$N^{-1/2} \exp\left[2\pi \mathrm{i} \left(\frac{k_a n_a}{a} + \frac{k_b n_b}{b} \right) \right].$$

The numerator $bk_a + ak_b$ is congruent to k modulo both a and b, hence modulo ab. This kernel equals $N^{-1/2} \exp(2\pi \mathrm{i} kn/N)$, proving (7). Omitting the dual-output permutation would generally give a different phase labeling.

1.3 What is lost when the sign is removed

The linear address-forgetting map

$$C_{\mathrm{phase}} |s, a\rangle = |a\rangle \tag{8}$$

maps a 2^{h+1}-dimensional dyadic unit register onto a 2^h-dimensional principal phase register. Its kernel is the span of the 2^h vectors $|0, a\rangle - |1, a\rangle$. Thus

$$\operatorname{rank} C_{\mathrm{phase}} = 2^h, \qquad \dim \ker C_{\mathrm{phase}} = 2^h.$$

This is a rank-deficient linear sum of amplitudes. It is not an isometry, and it is not the quantum channel that discards a subsystem. For example it sends the normalized state $(|0, a\rangle + |1, a\rangle)/\sqrt{2}$ to $\sqrt{2}|a\rangle$. Its adjoint product therefore does not equal the identity.

Discarding the sign as a quantum operation instead takes the partial trace $\rho \mapsto \mathrm{Tr}_{\mathrm{sign}}\, \rho$. Both pure states $|0, a\rangle$ and $|1, a\rangle$ then produce the same reduced phase state $|a\rangle\langle a|$. Yet the sign observable

$$Z_{\mathrm{sign}} = (|0\rangle\langle 0| - |1\rangle\langle 1|) \otimes I$$

has expectations $+1$ and -1 on those states. No measurement on the retained phase state can reconstruct this difference. The obstruction is information loss in a stated map, not an alternative probability rule.

A separate arithmetic example shows why integral residue selectors and rational spectral idempotents should be distinguished. Let

$$A = \begin{pmatrix} 1 & 1 \\ 0 & -1 \end{pmatrix}, \qquad A^2 = I, \qquad E_{\pm} = \frac{I \pm A}{2}.$$

The identities $E_{\pm}^2 = E_{\pm}$, $E_+E_- = 0$, and $E_+ + E_- = I$ follow directly from $A^2 = I$. Both idempotents contain a denominator two, and the resultant of $x - 1$ and $x + 1$ has dyadic valuation one. Modulo two the two roots agree and $(A - I)^2 = 0$, while $A - I$ is nonzero. Thus the residual matrix can retain a nilpotent extension even though it no longer distinguishes the two rational eigenvalues. These $E_{\pm}$ are complementary algebraic spectral idempotents. Since A is not self-adjoint in the standard counting inner product, they are not orthogonal quantum measurement projectors in that inner product. The unit-address sign projector, by contrast, is an ordinary diagonal orthogonal projector.

2 Cylinder averages, Haar details, and digit analysis

The principal phase register has a canonical finite filtration. For a prime p and depth $h \geq 0$, set

$$G_{p,h} = \mathbb{Z}/p^h\mathbb{Z}, \qquad \mathcal{H}_{p,h} = \ell^2(G_{p,h}; \mathbb{C}).$$

At depth zero this is a one-dimensional space. A level-r cylinder contains all residues with a prescribed prefix modulo p^r, and has p^{h-r} members. Throughout this section depth refers to the number of retained least significant base-p digits.

2.1 Orthogonal averaging and its detail resolution

For $0 \leq r \leq h$, define the cylinder average

$$(E_r^{(h)} f)(x) = p^{-(h-r)} \sum_{y \equiv x \pmod{p^r}} f(y). \tag{9}$$

Its matrix is real and symmetric. Averaging twice over the same cylinder has the same effect as averaging once. Thus $E_r^{(h)}$ is the orthogonal projector onto functions constant on level-r cylinders. The range has dimension p^r.

Proposition 2.1 (Finite Haar resolution). *The cylinder projectors satisfy*

$$E_r^{(h)} E_s^{(h)} = E_{\min(r,s)}^{(h)}, \qquad E_h^{(h)} = I.$$

Define

$$D_0^{(h)} = E_0^{(h)}, \qquad D_r^{(h)} = E_r^{(h)} - E_{r-1}^{(h)} \quad (1 \le r \le h).$$

Then the $D_r^{(h)}$ are mutually orthogonal projectors and

$$\begin{aligned} \sum_{r=0}^{h} D_r^{(h)} &= I, & D_r^{(h)} D_s^{(h)} &= \delta_{rs} D_r^{(h)}, \\ \operatorname{rank} D_0^{(h)} &= 1, & \operatorname{rank} D_r^{(h)} &= (p-1)p^{r-1} \quad (r \ge 1). \end{aligned} \tag{10}$$

Proof. If $r \le s$, the level-s cylinders partition each level-r cylinder into equally sized groups. Averaging over the smaller groups and then over the whole group returns the full average, giving $E_r E_s = E_r$. The self-adjointness of both sides also gives $E_s E_r = E_r$. This proves the minimum rule.

For $r \ge 1$, $(E_r - E_{r-1})^2 = E_r - 2E_{r-1} + E_{r-1} = D_r$. For $1 \le r < s$, expansion of $(E_r - E_{r-1})(E_s - E_{s-1})$ gives zero by the same minimum rule. The products with D_0 vanish as well. Telescoping gives the resolution of the identity. Finally $\operatorname{Ran} E_{r-1} \subset \operatorname{Ran} E_r$, so the detail rank is $p^r - p^{r-1}$, while E_0 has rank one. □

The decomposition is exact at each finite depth:

$$\mathcal{H}_{p,h} = \bigoplus_{r=0}^{h} \operatorname{Ran} D_r^{(h)}, \qquad \|f\|^2 = \sum_{r=0}^{h} \left\| D_r^{(h)} f \right\|^2. \tag{11}$$

A depth cutoff $m \le h$ retains $E_m^{(h)} = \sum_{r=0}^{m} D_r^{(h)}$. Its discarded norm is exactly

$$\left\| (I - E_m^{(h)}) f \right\|^2 = \sum_{r=m+1}^{h} \left\| D_r^{(h)} f \right\|^2.$$

No smoothness assumption is needed for this finite Parseval identity. A rate of decrease with m requires an additional bound on the detail weights of the state.

2.2 The child characters form the Haar basis

Write $\omega_p = \exp(2\pi \mathrm{i}/p)$. For $1 \le r \le h$, parent prefix $a \in \{0, \dots, p^{r-1}-1\}$, and character $q \in \{1, \dots, p-1\}$, define

$$\psi_{r,a,q}^{(h)}(x) = \begin{cases} p^{-(h-r+1)/2} \omega_p^{q d_r(x)}, & x \equiv a \pmod{p^{r-1}}, \\ 0, & \text{otherwise}, \end{cases} \tag{12}$$

where $d_r(x)$ is the digit multiplying p^{r-1}. Include the constant scaling vector $\psi_0^{(h)}(x) = p^{-h/2}$.

Each wavelet is constant on level-r cylinders and has zero average on its parent level-$r-1$ cylinder, because $\sum_{d=0}^{p-1} \omega_p^{qd} = 0$ for nonzero q. Thus it lies in $\operatorname{Ran} D_r^{(h)}$. Its support contains p^{h-r+1} points, giving

unit norm. Distinct parents have disjoint supports. For a fixed parent the inner product of the characters q, q' is

$$p^{-1}\sum_{d=0}^{p-1}\omega_p^{(q'-q)d}=\delta_{qq'}.$$

Wavelets at different levels are orthogonal because their detail projectors are orthogonal. There are exactly $(p-1)p^{r-1}$ wavelets at level r, equal to the rank in (10), so they span that detail subspace. Hence

$$D_r^{(h)}=\sum_{a=0}^{p^{r-1}-1}\sum_{q=1}^{p-1}|\psi_{r,a,q}^{(h)}\rangle\langle\psi_{r,a,q}^{(h)}|. \tag{13}$$

The omitted character $q=0$ is the parent average. For $p=2$ there is one average and one difference per parent. For odd p there is one average and $p-1$ nontrivial child characters. The dyadic sign factor of (3) is independent of this average/detail split in the principal phase register.

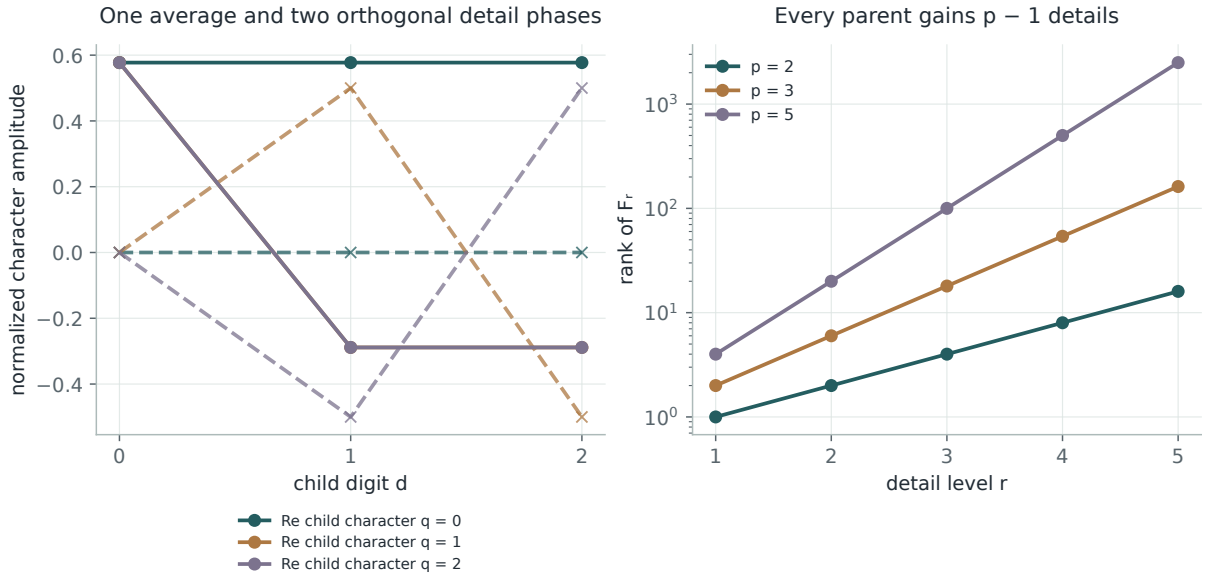


Figure 1. For $p=3$, the normalized child characters contain one average and two orthogonal complex details. Solid curves show real parts and dashed curves show imaginary parts. The level rank $(p-1)p^{r-1}$ grows by the branching factor. The atlas uses F_r for the corresponding infinite cylinder detail; its finite restriction is $D_r^{(h)}$ for $r\le h$.

2.3 The refinement isometry and its adjoint

One more digit supplies p children for every address. The uniform refinement map is

$$(V_{p,h}f)(a+dp^h)=p^{-1/2}f(a),\quad 0\le a<p^h,\quad 0\le d<p. \tag{14}$$

In tensor coordinates it is $f\mapsto|+_p\rangle\otimes f$, where $|+_p\rangle=p^{-1/2}\sum_d|d\rangle$ and the child digit is listed first. The norm calculation is

$$\|V_{p,h}f\|^2=\sum_a\sum_d p^{-1}|f(a)|^2=\|f\|^2.$$

The adjoint is the normalized amplitude sum

$$(V_{p,h}^{*}g)(a) = p^{-1/2}\sum_{d=0}^{p-1} g(a+dp^{h}). \tag{15}$$

Consequently

$$V_{p,h}^{*}V_{p,h} = I, \qquad V_{p,h}V_{p,h}^{*} = E_{h}^{(h+1)}, \qquad I - V_{p,h}V_{p,h}^{*} = D_{h+1}^{(h+1)}. \tag{16}$$

The first composition is an identity on the coarse space. The second is a projector on the fine space, not a fine-space identity.

Replicated vectors are constant in the new digit. All old cylinder averages commute with replication, giving

$$\begin{aligned} E_{r}^{(h+1)}V_{p,h} &= V_{p,h}E_{r}^{(h)} \quad (0 \le r \le h), \\ D_{r}^{(h+1)}V_{p,h} &= V_{p,h}D_{r}^{(h)} \quad (0 \le r \le h), \qquad\qquad D_{h+1}^{(h+1)}V_{p,h} = 0. \end{aligned} \tag{17}$$

Repeated refinements compose as isometries. For $w \ge 0$ and one coarse address,

$$V_{p,h+w-1}\cdots V_{p,h}|a\rangle = p^{-w/2}\sum_{\mathbf{d}\in\{0,\dots,p-1\}^{w}} |a,\mathbf{d}\rangle, \tag{18}$$

with the empty product interpreted as the identity. The factor $p^{-w/2}$ is forced by the isometry and orthonormality of the descendant addresses.

2.4 Digit reversal and the polyphase split

For $h \ge 1$, put $N = p^{h}$ and $m = p^{h-1}$. Two permutation unitaries separate different digits:

$$\begin{aligned} \mathcal{A}_{p,h}^{\mathrm{std}}|pk+r\rangle &= |r\rangle\otimes|k\rangle, & 0 \le r < p, \quad 0 \le k < m, \\ \mathcal{A}_{p,h}^{\mathrm{ad}}|a+dp^{h-1}\rangle &= |d\rangle\otimes|a\rangle, & 0 \le d < p, \quad 0 \le a < m. \end{aligned} \tag{19}$$

The standard polyphase split extracts the least significant residue. The adic split extracts the newest, most significant digit at the current depth. Let $\mathcal{R}_{p,h}$ reverse the h base-p digits, with $\mathcal{R}_{p,0} = I$ on the one-dimensional space. Then

$$\mathcal{A}_{p,h}^{\mathrm{ad}} = (I_{p}\otimes\mathcal{R}_{p,h-1})\mathcal{A}_{p,h}^{\mathrm{std}}\mathcal{R}_{p,h}. \tag{20}$$

Indeed reversing $d_{0}+d_{1}p+\cdots+d_{h-1}p^{h-1}$ makes d_{h-1} the least significant digit. Standard analysis removes that digit, and reversing the remaining $h-1$ digits restores the prefix $a = d_{0}+\cdots+d_{h-2}p^{h-2}$. This proves the identity on an orthonormal basis.

Let the new-digit modulation be $M_{p,h}|a+dp^{h-1}\rangle = \omega_{p}^{d}|a+dp^{h-1}\rangle$. The selector of child digit d is the exact roots-of-unity filter

$$R_{d} = \frac{1}{p}\sum_{\ell=0}^{p-1}\omega_{p}^{-d\ell}M_{p,h}^{\ell}. \tag{21}$$

On a child with digit e, the coefficient is $p^{-1}\sum_{\ell}\omega_{p}^{(e-d)\ell} = \delta_{de}$. Therefore $R_{d}R_{e} = \delta_{de}R_{d}$, $R_{d}^{*} = R_{d}$, and $\sum_{d}R_{d} = I$. These projectors select address digits. They are different from the projectors onto Fourier characters of a child register.

Apply the child Fourier transform after adic analysis:

$$\mathcal{W}_{p,h} = (F_p \otimes I_m)\mathcal{A}^{\mathrm{ad}}_{p,h}. \tag{22}$$

The zero child-frequency row of F_p is $p^{-1/2}(1, \dots, 1)$. Hence

$$\begin{aligned} (\langle 0| \otimes I_m)\mathcal{W}_{p,h} &= V^*_{p,h-1}, \\ \mathcal{W}^*_{p,h}[(I_p - |0\rangle\langle 0|) \otimes I_m]\mathcal{W}_{p,h} &= D^{(h)}_h. \end{aligned} \tag{23}$$

The second identity subtracts the coarse zero-frequency projector from the identity using (16). Recursive application to the zero-frequency channel alone constructs a unitary Haar analysis map $\mathcal{T}_{p,h}$ whose channel projectors Π_r have the ranks in (10) and obey

$$\mathcal{T}^*_{p,h}\Pi_r\mathcal{T}_{p,h} = D^{(h)}_r. \tag{24}$$

With our positive Fourier sign, the analysis row of character q pairs with the complex-conjugate wavelet character. This only permutes the nonzero character labels $q \leftrightarrow p - q$ and leaves the detail projectors unchanged.

The zero child-frequency channel retains a uniform superposition over the new digit. Retaining standard residue zero instead keeps one address out of each p addresses. Both spaces have the same dimension, but they are distinct subspaces. Their compressed dynamics and Schur complements consequently need not agree. Digit reversal and child Fourier analysis specify the transformations needed to compare them.

3 Which Hamiltonians are compatible with refinement?

An isometric address refinement does not automatically preserve a Hamiltonian. The exact condition is a block condition on that Hamiltonian. It can be checked without an approximation, a scaling limit, or a physical interpretation of the new digit.

Theorem 3.1 (Finite refinement classification)**.** *Let $V : \mathcal{H}_c \to \mathcal{H}_f$ be an isometry between finite Hilbert spaces. Set $P = VV^*$ and $Q = I - P$. For self-adjoint operators H_c, H_f, the following conditions are equivalent:*

$$H_f V = V H_c, \tag{25}$$

$$QH_f P = 0, \qquad PH_f P = V H_c V^*. \tag{26}$$

They hold exactly when, on $\operatorname{Ran} P \oplus \operatorname{Ran} Q$*,*

$$H_f = (V H_c V^*)|_{\operatorname{Ran} P} \oplus H_Q, \qquad H_Q = (QH_f Q)|_{\operatorname{Ran} Q}. \tag{27}$$

The detail operator H_Q may be any self-adjoint operator on $\operatorname{Ran} Q$*.*

Proof. Right multiplication of (25) by V^* gives $H_f P = V H_c V^*$. Left multiplication by Q and P gives (26). Conversely $PV = V$ and $V^*V = I$ imply

$$H_f V = (PH_f P + QH_f P)V = V H_c V^* V = V H_c.$$

Self-adjointness gives $PH_fQ = (QH_fP)^* = 0$. Thus both off-diagonal blocks vanish, and the remaining blocks are exactly those in (27). Conversely that direct sum has the two block properties. There is no further condition on the self-adjoint detail block. □

The distinction between invariance and compression is essential. Defining $H_c = V^*H_fV$ always matches the coarse block, but leaves QH_fV unconstrained. Thus compression alone does not establish (25). When it does hold, the coarse subspace is reducing, and fine evolution introduces no new detail amplitude from a coarse initial state.

Corollary 3.2 (Functional-calculus intertwining)**.** *Under Theorem 3.1, for every complex function f defined on the finite union of the spectra,*

$$f(H_f)V = Vf(H_c).$$

In particular, for real t, real β, and $z \notin \sigma(H_f)$,

$$\begin{aligned} e^{-\mathrm{i}tH_f}V &= Ve^{-\mathrm{i}tH_c}, \\ e^{-\beta H_f}V &= Ve^{-\beta H_c}, \\ (z - H_f)^{-1}V &= V(z - H_c)^{-1}. \end{aligned} \tag{28}$$

Proof. On Ran P, V is a unitary identification of H_c with the first block of H_f. The finite spectral decomposition therefore maps each coarse spectral subspace to the corresponding first-block spectral subspace. Applying f to the eigenvalues gives the identity. The three displayed functions are defined on the relevant spectra. □

The Euclidean operator $e^{-\beta H}$ in this corollary is a spectral multiplier. It is not generally unitary, and a normalized state after Euclidean evolution requires its own normalization. The identity does not compare normalized Gibbs states of the entire fine and coarse systems, since the fine partition function also counts the detail block. For a compatible finite tower, the coarse propagator and its embedded fine propagator agree exactly. This does not make an arbitrary interacting fine Hamiltonian compatible.

3.1 Hierarchical and translation-invariant examples

Choose a real sequence of level energies $\epsilon_{p,r}$ and define

$$H_{p,h}^{\mathrm{Haar}} = \sum_{r=0}^{h} \epsilon_{p,r} D_r^{(h)}. \tag{29}$$

Every wavelet in Ran $D_r^{(h)}$ has energy $\epsilon_{p,r}$. The projectors in (17) show directly that $H_{p,h+1}^{\mathrm{Haar}} V_{p,h} = V_{p,h} H_{p,h}^{\mathrm{Haar}}$. The new detail energy $\epsilon_{p,h+1}$ does not alter the embedded older levels. Moreover the recursive Haar transform diagonalizes the level blocks:

$$\mathcal{T}_{p,h} H_{p,h}^{\mathrm{Haar}} \mathcal{T}_{p,h}^* = \sum_{r=0}^{h} \epsilon_{p,r} \Pi_r.$$

This is exact for each finite h and every finite real choice of the energies. Their growth becomes a domain question only when an infinite-depth operator is constructed.

Let the cyclic translation be $(S_{p,h}f)(x) = f(x-1)$, with subtraction modulo p^h. Since a fine replicated function depends only on the residue modulo p^h,

$$S_{p,h+1}V_{p,h} = V_{p,h}S_{p,h}.$$

The adjoint translation obeys the same identity. Therefore the cyclic Laplacian $L_{p,h} = 2I - S_{p,h} - S_{p,h}^*$ intertwines, and for real κ_p the family

$$H_{p,h} = H_{p,h}^{\mathrm{Haar}} + \kappa_p L_{p,h} \tag{30}$$

is self-adjoint and refinement compatible. Translation permutes the level-r cylinders, so it commutes with each $E_r^{(h)}$ and $D_r^{(h)}$. The Laplacian thus preserves every detail subspace, although a chosen spatially localized Haar wavelet need not be its eigenvector.

This cyclic translation uses fixed unit displacement in the address group. Multiplying it by a depth-dependent factor intended as a physical inverse lattice spacing changes the tower. Such a factor is not covered by (30) unless the resulting operators still satisfy the stated intertwining equation. Exact adic compatibility and continuum spatial consistency are different requirements.

3.2 A valuation-shell propagator

For the pure hierarchical Hamiltonian, the finite spectral theorem gives

$$U_{p,h}(t) = e^{-\mathrm{i}tH_{p,h}^{\mathrm{Haar}}} = \sum_{r=0}^{h} e^{-\mathrm{i}t\epsilon_{p,r}} D_r^{(h)}. \tag{31}$$

The cylinder kernel is

$$E_r^{(h)}(x,y) = p^{-(h-r)}\mathbf{1}_{\{x\equiv y \,(\mathrm{mod}\, p^r)\}}.$$

Consequently

$$\begin{aligned} U_{p,h}(t;x,y) &= p^{-h}e^{-\mathrm{i}t\epsilon_{p,0}} \\ &\quad + \sum_{r=1}^{h} e^{-\mathrm{i}t\epsilon_{p,r}} \left[p^{-(h-r)}\mathbf{1}_{x\equiv y\,(\mathrm{mod}\,p^r)} - p^{-(h-r+1)}\mathbf{1}_{x\equiv y\,(\mathrm{mod}\,p^{r-1})}\right]. \end{aligned} \tag{32}$$

Every indicator is determined by the truncated valuation $\min\{v_p(x-y), h\}$, with the diagonal assigned value h. The transition amplitude, and therefore the individual-address Born probability, is constant on each such valuation shell. At a fixed initial address y, the shell of valuation $k < h$ contains $(p-1)p^{h-k-1}$ final addresses. Its total probability includes this multiplicity. The diagonal has one address. Uniform probability per address on a shell does not mean equal total probabilities of different shells.

3.3 Leakage detects failure of compatibility

For a self-adjoint fine Hamiltonian with coarse projector P, set $B = QH_fP$. If a unit vector ψ lies in Ran P, its initial detail amplitude is

$$Qe^{-\mathrm{i}tH_f}\psi = -\mathrm{i}tB\psi + O(t^2), \qquad \left\|Qe^{-\mathrm{i}tH_f}\psi\right\|^2 = t^2\left\|B\psi\right\|^2 + O(|t|^3).$$

These are finite Taylor expansions for fixed H_f and ψ. Thus a nonzero B supplies a state with nonzero short-time leakage. The two-dimensional example

$$H_f = \begin{pmatrix} 0 & \overline{\gamma} \\ \gamma & 0 \end{pmatrix}, \qquad P = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$$

has compressed Hamiltonian zero but detail probability $\sin^2(|\gamma|t)$ from the first basis state. It fails exact refinement compatibility whenever $\gamma \neq 0$.

For z outside the spectrum of $H_Q = (QH_fQ)|_{\mathrm{Ran}\,Q}$, the coarse Schur self-energy is

$$\Sigma_P(z) = PH_fQ(z - H_Q)^{-1}QH_fP. \tag{33}$$

Compatibility makes it identically zero by setting $B = 0$. This zero is an algebraic statement, not a uniform bound on a resolvent near its spectrum. A nonzero coupling requires the energy-dependent Schur term when the detail sector is eliminated. The spectral part of this monograph develops that elimination and its time-domain memory.

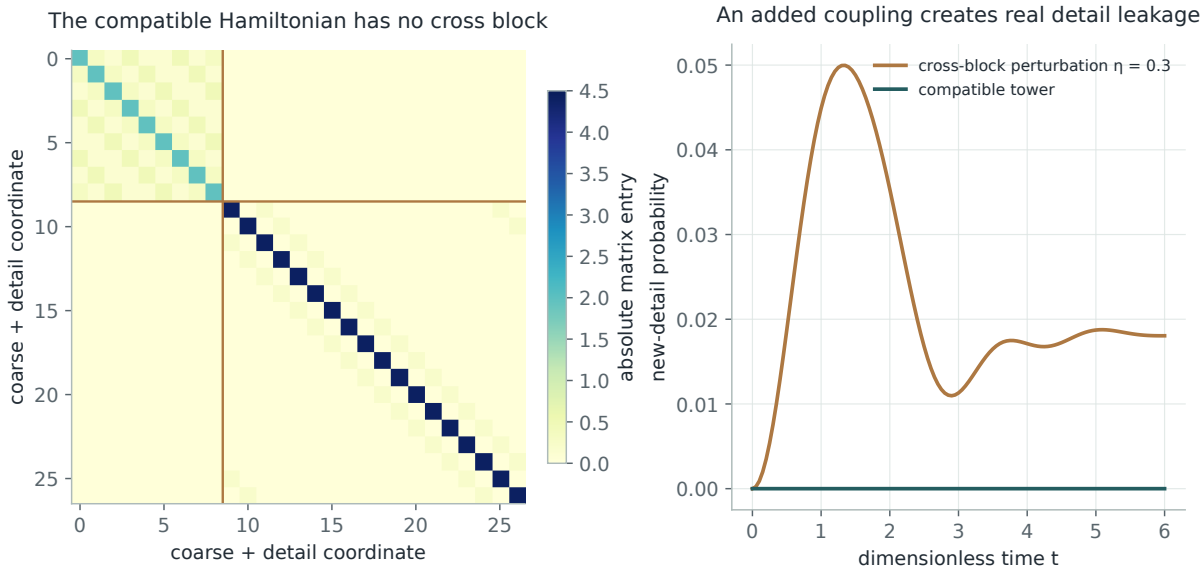


Figure 2. A finite compatible Hamiltonian has zero blocks between its embedded coarse space and its new detail space. In the plotted example, adding an off-diagonal coupling creates detail probability from a coarse initial state. This finite example illustrates the block criterion and supplies a negative control for unconditional refinement claims.

4 Normalization and collective coupling

Refinement fixes the norm of an embedded packet. It does not fix an interaction matrix element. The distinction becomes visible when one coarse packet couples to many descendants rather than to one already normalized aggregate packet. In this section $\hbar = 1$, so the Hamiltonian has units of inverse time.

4.1 Descendant weights are fixed by the isometry

After w refinements a packet has $n = p^w$ orthonormal descendants. Its normalized radial vector is

$$|R_{p,w}\rangle = n^{-1/2} \sum_{\mathbf{d}} |\mathbf{d}\rangle. \tag{34}$$

The norm is $n^{-1} \sum_{\mathbf{d}} 1 = 1$, so each descendant has amplitude modulus $p^{-w/2}$ and Born probability p^{-w}. This reproduces (18). The words "radial" and "angular" here refer to the uniform vector and its orthogonal complement in a finite descendant space. They do not introduce a spacetime radius or a rotation group.

Introduce a visible address $|c\rangle$ orthogonal to all descendants. Suppose every descendant has the same coupling from $|c\rangle$, with complex matrix element $\gamma_{p,w}$. The Hermitian interaction is

$$H_{\text{int}} = \sum_{\mathbf{d}} \left(\gamma_{p,w} |\mathbf{d}\rangle\langle c| + \overline{\gamma_{p,w}} |c\rangle\langle \mathbf{d}| \right). \tag{35}$$

The conjugate reverse coefficient is necessary when $\gamma_{p,w}$ is not real. Define the collective coefficient

$$G_{p,w} = \sqrt{n}\, \gamma_{p,w}. \tag{36}$$

Then $H_{\text{int}}|c\rangle = G_{p,w}|R_{p,w}\rangle$ and $H_{\text{int}}|R_{p,w}\rangle = \overline{G_{p,w}}|c\rangle$. If $|A\rangle = \sum_{\mathbf{d}} a_{\mathbf{d}} |\mathbf{d}\rangle$ is orthogonal to $|R_{p,w}\rangle$, its coefficients satisfy $\sum_{\mathbf{d}} a_{\mathbf{d}} = 0$. Hence $H_{\text{int}}|A\rangle = 0$. Only the radial vector is bright under this equal-leaf interaction.

This statement concerns the interaction block. A separate descendant Hamiltonian can mix radial and angular vectors. To obtain an invariant two-dimensional radial subsystem, the full descendant Hamiltonian must preserve the radial line, and the full Hamiltonian must leave its angular complement uncoupled from $|c\rangle$. A degenerate hidden energy ϵI is one sufficient choice.

4.2 The exact bright two-level problem

Assume those invariance conditions and order the basis as $(|c\rangle, |R_{p,w}\rangle)$. With visible energy E_c and radial energy ϵ, the exact bright block is

$$H_{\text{bright}} = \begin{pmatrix} E_c & \overline{G}_{p,w} \\ G_{p,w} & \epsilon \end{pmatrix}. \tag{37}$$

Set $\overline{E} = (E_c + \epsilon)/2$, $\delta = (E_c - \epsilon)/2$, and $\Omega = (\delta^2 + |G_{p,w}|^2)^{1/2}$. Then $(H_{\text{bright}} - \overline{E} I)^2 = \Omega^2 I$, giving

$$e^{-\mathrm{i}tH_{\text{bright}}} = e^{-\mathrm{i}t\overline{E}} \left[I \cos(\Omega t) - \mathrm{i} \frac{H_{\text{bright}} - \overline{E} I}{\Omega} \sin(\Omega t) \right]. \tag{38}$$

At $\Omega = 0$ this formula is understood by continuous extension. Starting in $|c\rangle$, the radial probability is therefore

$$\Pr(R_{p,w}; t) = \frac{|G_{p,w}|^2}{\delta^2 + |G_{p,w}|^2} \sin^2\left(t \sqrt{\delta^2 + |G_{p,w}|^2} \right). \tag{39}$$

The probability is defined to be zero when $\Omega = 0$, equivalently when both δ and $G_{p,w}$ vanish. On resonance $E_c = \epsilon$ it reduces to $\sin^2(|G_{p,w}|t)$. The individual descendant probabilities are $\Pr(R_{p,w}; t)/p^w$ when the hidden amplitude stays radial.

Eliminating the hidden radial level gives

$$\Sigma_{p,w}(z) = \frac{|G_{p,w}|^2}{z - \epsilon}, \qquad z \neq \epsilon. \tag{40}$$

This is the exact finite Schur term. The full visible Green function, when its denominator is nonzero, is $[z - E_c - \Sigma_{p,w}(z)]^{-1}$. It is distinct from the self-energy itself, and the locations of its poles depend on both sectors.

4.3 The coupling exponent is a model choice

Let $g \in \mathbb{C}$ be fixed and choose $\gamma_{p,w} = gp^{-\alpha w}$ for a real exponent α. Then

$$G_{p,w} = gp^{(1/2-\alpha)w}, \qquad |G_{p,w}|^2 = |g|^2 p^{(1-2\alpha)w}. \tag{41}$$

The exact branching multiplicity produces the factor $p^{w/2}$. Thus the collective coefficient grows with w for $\alpha < 1/2$, is independent of w for $\alpha = 1/2$, and decays for $\alpha > 1/2$. The same cases govern the operator norm of this equal-leaf interaction: its two radial eigenvalues are $\pm|G_{p,w}|$, so $\|H_{\text{int}}\| = |G_{p,w}|$. They do not determine the boundedness of a different infinite operator with many interacting depth sectors.

Per-leaf choice	α	$\lvert G_{p,w}\rvert$	$\lvert G_{p,w}\rvert^2$
Constant matrix element	0	$\lvert g\rvert p^{w/2}$	$\lvert g\rvert^2 p^w$
Constant collective coupling	$1/2$	$\lvert g\rvert$	$\lvert g\rvert^2$
Uniform measure coefficient	1	$\lvert g\rvert p^{-w/2}$	$\lvert g\rvert^2 p^{-w}$

The measure coefficient $\alpha = 1$ assigns the number p^{-w} to each per-leaf matrix element after factoring out g. It gives a collective amplitude of size $p^{-w/2}$ and a self-energy numerator of size p^{-w}. The self-energy estimate is uniform only where $|z - \epsilon|$ stays bounded below. Likewise

$$\Pr(R_{p,w}; t) \leq |g|^2 t^2 p^{-w} \quad \text{when } \alpha = 1$$

follows from (39) and $|\sin u| \leq |u|$. It is a fixed-time bound. On resonance a time of order $p^{w/2}/|g|$ can still transfer order-one probability.

A model that directly couples $|c\rangle$ to one normalized hidden packet with matrix element gp^{-w} instead has a collective coefficient of that same size and a self-energy numerator $|g|^2 p^{-2w}$. It is a valid two-level model, but it is not the equal-leaf model (35). The number of descendant addresses alone does not select between these interaction geometries. A physical refinement law must specify the coupling, the retained hidden modes, their energies, and the observation scale.

Radial elimination and polyphase spectra

5 Polyphase coordinates and folded bands

The finite Haar decomposition isolates an average and its new detail channels. A standard polyphase decomposition instead groups the entries of a cyclic lattice by their residue class. Both transformations preserve the full state space. An elimination begins only when some coordinates are solved for and removed. The spectral parameter then enters the remaining operator, and the eliminated degrees of freedom return as memory in time.

We use complex Hilbert spaces with inner product linear in the second argument and set $\hbar = 1$. All matrices in the finite sections act on their stated coordinate spaces. In particular, the inverse of a low channel block is a low channel operator, rather than a full space matrix.

5.1 Residue coordinates and exact block reconstruction

Fix an integer $p \geq 2$ and $M \geq 3$, and put $N = pM$. Primality is not needed for the identities in this section. For $0 \leq r < p$, define the isometry $J_r : \mathbb{C}^M \to \mathbb{C}^N$ by

$$J_r|k\rangle = |pk + r\rangle, \qquad 0 \leq k < M.$$

Then $J_r^* J_s = \delta_{rs} I_M$ and $\sum_r J_r J_r^* = I_N$. The standard analysis map $\mathcal{A}$ sends $|pk + r\rangle$ to $|r\rangle \otimes |k\rangle$. This is a permutation unitary. For any operator X on $\mathbb{C}^N$, its residue blocks $X_{rs} = J_r^* X J_s$ reconstruct it:

$$X = \sum_{r,s=0}^{p-1} J_r X_{rs} J_s^*, \qquad (\mathcal{A} X \mathcal{A}^*)_{rs} = X_{rs}. \tag{42}$$

The first equality follows by inserting $\sum_r J_r J_r^*$ on both sides of X. If $H = H^*$, its blocks satisfy $H_{rs}^* = H_{sr}$. If U is unitary, the same insertion gives

$$\begin{aligned} \sum_r U_{rs}^* U_{rt} &= J_s^* U^* U J_t = \delta_{st} I_M, \\ \sum_s U_{rs} U_{ts}^* &= J_r^* U U^* J_t = \delta_{rt} I_M. \end{aligned} \tag{43}$$

These are the block unitary identities. Every coordinate is retained.

For finite H, expansion of the matrix exponential proves

$$e^{-it\mathcal{A}H\mathcal{A}^*} = \mathcal{A} e^{-itH} \mathcal{A}^*.$$

Indeed $(\mathcal{A}H\mathcal{A}^*)^n = \mathcal{A}H^n\mathcal{A}^*$, and the series converges in matrix norm. The same transformation of each Kraus operator gives a unitarily equivalent quantum channel: if $\Phi(\rho) = \sum_\mu K_\mu \rho K_\mu^*$ with $\sum_\mu K_\mu^* K_\mu = I$, then

$$K_\mu^{\mathrm{poly}} = \mathcal{A} K_\mu \mathcal{A}^*, \qquad \sum_\mu (K_\mu^{\mathrm{poly}})^* K_\mu^{\mathrm{poly}} = I.$$

This is a change of coordinates in a given channel. It does not supply a new dissipative dynamics.

Proposition 5.1 (Probability currents between residue channels). *Suppose $i\dot{\psi}_r = \sum_s H_{rs}\psi_s$ with $H^*_{rs} = H_{sr}$. The channel weights $w_r = \|\psi_r\|^2$ satisfy*

$$\dot{w}_r = \sum_s \mathcal{J}_{r\leftarrow s}, \qquad \mathcal{J}_{r\leftarrow s} = 2\,\mathrm{Im}\langle\psi_r, H_{rs}\psi_s\rangle, \qquad \mathcal{J}_{r\leftarrow s} = -\mathcal{J}_{s\leftarrow r}.$$

Consequently $\sum_r w_r$ is constant.

Proof. Differentiating the norm gives $\dot{w}_r = 2\,\mathrm{Re}\langle\psi_r, -i\sum_s H_{rs}\psi_s\rangle$, which is the stated sum of imaginary parts. Self-adjoint block symmetry implies $\langle\psi_s, H_{sr}\psi_r\rangle = \overline{\langle\psi_r, H_{rs}\psi_s\rangle}$. The currents therefore cancel in pairs, including a zero diagonal current. □

Here w_r is a probability or squared amplitude weight. Calling it an energy would require an additional energy observable and its balance law.

5.2 A Bloch symbol with a fixed shift convention

Let $S|n\rangle = |n+1 \pmod N\rangle$ and

$$H = \mu I + \kappa(S + S^*), \qquad \mu, \kappa \in \mathbb{R}. \tag{44}$$

Use the coarse Fourier matrix $(F_M)_{jk} = M^{-1/2}e^{-2\pi ijk/M}$ and $\theta_j = 2\pi j/M$. After $\mathcal{A}$, Fourier transformation in k, and the permutation from residue major to momentum major ordering, the operator becomes a direct sum of $p \times p$ symbols. To see their entries, note that S sends (r,k) to $(r+1,k)$ for $r < p-1$ and sends $(p-1,k)$ to $(0,k+1)$. The coarse shift has Fourier eigenvalue $e^{-i\theta_j}$. Thus define

$$P_p(\theta)e_r = \begin{cases} e_{r+1}, & r < p-1, \\ e^{-i\theta}e_0, & r = p-1. \end{cases} \qquad \mathsf{H}_p(\theta) = \mu I_p + \kappa(P_p(\theta) + P_p(\theta)^*). \tag{45}$$

The transformed H is $\bigoplus_{j=0}^{M-1} \mathsf{H}_p(\theta_j)$. In particular, $P_p(\theta)^p = e^{-i\theta}I_p$.

Proposition 5.2 (Exact folded cyclic spectrum). *For $q = 0, \dots, p-1$, the eigenvalues of $\mathsf{H}_p(\theta)$ are*

$$E_q(\theta) = \mu + 2\kappa\cos\left(\frac{\theta + 2\pi q}{p}\right). \tag{46}$$

At $\theta = \theta_j$, the union of these pM eigenvalues, counting multiplicity, is exactly the spectrum of (44).

Proof. Put $\alpha_q = (\theta + 2\pi q)/p$ and let the rth component of v_q be $p^{-1/2}e^{i\alpha_q r}$. On the interior components, $P_p(\theta)v_q = e^{-i\alpha_q}v_q$. The wrapped component has the same value because $e^{ip\alpha_q} = e^{i\theta}$. The v_q are orthonormal by the finite character sum. Applying $\mu I + \kappa(P_p + P_p^*)$ gives (46). Every fine Fourier index is uniquely $k = j + qM$. Its angle satisfies $2\pi k/(pM) = (\theta_j + 2\pi q)/p$. Substitution into the fine cyclic eigenvalue $\mu + 2\kappa\cos(2\pi k/N)$ proves the multiplicity statement. □

The same construction works for any circulant polynomial $H = \sum_\ell c_\ell S^\ell$, with symbol $\sum_\ell c_\ell P_p(\theta)^\ell$. It preserves all bands, whether or not their energies coincide. For $N = p^h$, digit reversal relates this standard residue split to the newest-digit adic split developed earlier. It is still necessary to distinguish the retained residue-zero sites from the retained Haar average channel. They are different subspaces.

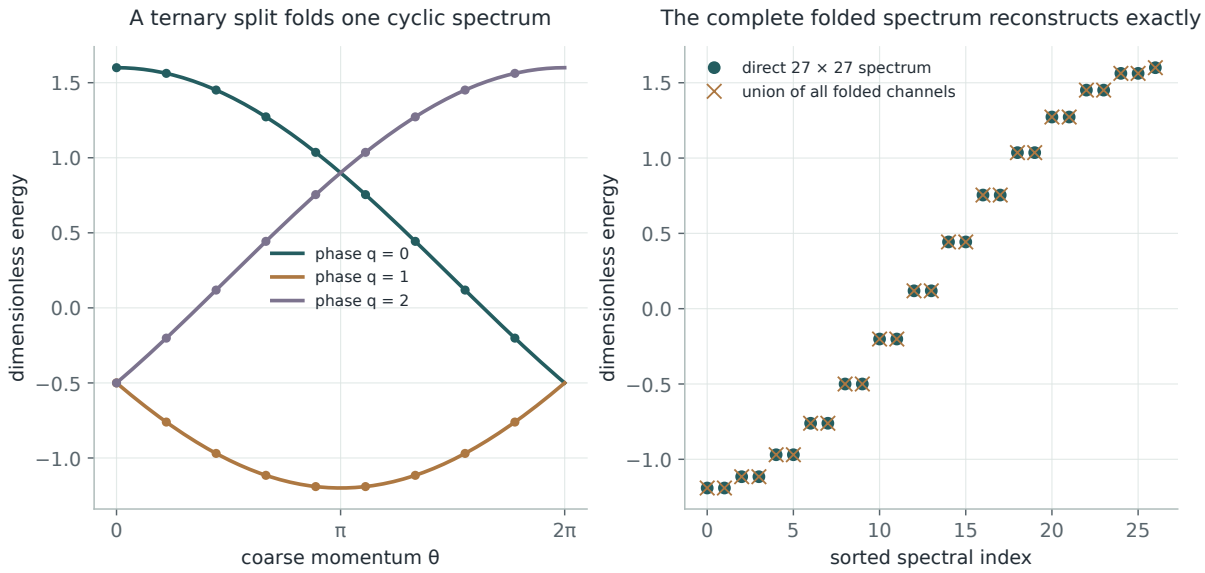


Figure 3. Three exact folded bands for the nearest-neighbor cycle. The discrete markers represent the finite cyclic spectrum. Band folding reorganizes its multiplicities without discarding any eigenstate.

6 Regular Schur reduction and causal memory

Let a finite Hilbert space be the orthogonal sum $\mathcal{H}_L \oplus \mathcal{H}_D$, and write a self-adjoint Hamiltonian as

$$H = \begin{pmatrix} A & B \\ B^* & D \end{pmatrix}, \qquad A = A^*, \quad D = D^*. \tag{47}$$

Here $B : \mathcal{H}_D \to \mathcal{H}_L$. Let ι_L embed $\mathcal{H}_L$ into the full space. This notation fixes the coordinate type of every compression.

6.1 The resolvent chart and its exceptional energies

Theorem 6.1 (Finite Feshbach identity). *If $z \notin \sigma(D)$, define*

$$\Sigma(z) = B(z - D)^{-1}B^*, \qquad \mathcal{S}(z) = z - A - \Sigma(z). \tag{48}$$

Then $z - H$ is invertible exactly when $\mathcal{S}(z)$ is invertible. On this common regular set,

$$\begin{aligned} \iota_L^*(z - H)^{-1}\iota_L &= \mathcal{S}(z)^{-1}, \\ \det(z - H) &= \det(z - D)\det \mathcal{S}(z). \end{aligned} \tag{49}$$

For $E \in \mathbb{R} \setminus \sigma(D)$, an eigenvector of H is equivalent to a nonzero solution of $\mathcal{S}(E)u = 0$, with hidden component $v = (E - D)^{-1}B^*u$.

Proof. Put $C = z - D$. Direct block multiplication gives

$$z - H = \begin{pmatrix} I & -BC^{-1} \\ 0 & I \end{pmatrix} \begin{pmatrix} \mathcal{S}(z) & 0 \\ 0 & C \end{pmatrix} \begin{pmatrix} I & 0 \\ -C^{-1}B^* & I \end{pmatrix}.$$

Both outer factors are invertible and have determinant one. This proves the invertibility equivalence and determinant formula. Inverting the factors gives the stated upper left block. Alternatively, in $(z - H)(u, v)^{\mathsf{T}} = (f, 0)^{\mathsf{T}}$, the lower equation gives $v = C^{-1}B^{*}u$ and the upper equation becomes $\mathcal{S}(z)u = f$. The same elimination with $f = 0$ proves the eigenvector statement. If $u = 0$ at such an E, invertibility of $E - D$ forces $v = 0$, so a nonzero full eigenvector indeed has $u \neq 0$. □

If $P_L = \iota_L \iota_L^{*}$, the full space compression is $P_L(z - H)^{-1}P_L = \iota_L \mathcal{S}(z)^{-1}\iota_L^{*}$. When a Haar transform $\mathcal{W}$ first exposes low and detail channels, the embedding in the original coordinates is $\iota_L = \mathcal{W}^{*}(|0\rangle \otimes I)$. The inverse on $\mathcal{H}_L$ must be embedded in this way before it is compared with a full space matrix.

The determinant equality is initially a rational identity on the regular set. Multiplication by $\det(z-D)$ cancels its denominators and yields the full characteristic polynomial, which extends across $\sigma(D)$. This extension does not make the displayed inverse of $z - D$ exist at a pole. At a resonant energy the uneliminated equations determine which hidden modes couple, survive, or cancel. A resonant subspace can be kept as part of an enlarged visible block.

6.2 Spectral residues, analytic sign, and branch weights

Let $D = \sum_{\lambda} \lambda E_{\lambda}$ be its spectral resolution. Then

$$\Sigma(z) = \sum_{\lambda} \frac{BE_{\lambda}B^{*}}{z - \lambda}. \tag{50}$$

Every residue $BE_{\lambda}B^{*}$ is positive semidefinite. Its rank records which directions of the hidden eigenspace reach the visible space. A vanishing residue removes the corresponding pole of Σ, even though $z - D$ is still singular there.

Proposition 6.2 (Upper half-plane sign and the exact visible norm)**.** *For $z = x + iy$ with $y > 0$,*

$$\operatorname{Im}\Sigma(z) = -yB\big((x - D)^2 + y^2 I\big)^{-1}B^{*} \leq 0, \qquad \operatorname{Im}\mathcal{S}(z) \geq yI.$$

In particular $\mathcal{S}(z)$ is invertible and $\|\mathcal{S}(z)^{-1}\| \leq y^{-1}$. For real $E \notin \sigma(D)$,

$$\Sigma'(E) = -B(E - D)^{-2}B^{*}.$$

If (u, v) is a normalized full eigenvector at E, then

$$1 = \langle u, (I - \Sigma'(E))u\rangle. \tag{51}$$

Proof. For real λ, $\operatorname{Im}(z - \lambda)^{-1} = -y/((x - \lambda)^2 + y^2)$. Applying this identity to the spectral sum gives the first formula and then the bound on $\operatorname{Im}\mathcal{S}$. For any u, Cauchy–Schwarz implies $y\|u\|^2 \leq |\langle u, \mathcal{S}(z)u\rangle| \leq \|u\|\,\|\mathcal{S}(z)u\|$. Thus $\mathcal{S}(z)$ has trivial kernel and an inverse of the stated norm in finite dimension. Differentiating the resolvent gives the minus sign in Σ'. The reconstructed hidden norm is $\|v\|^2 = \langle u, B(E - D)^{-2}B^{*}u\rangle$. Adding $\|u\|^2$ proves (51). □

For a one-dimensional visible block this says that an eigenstate's visible probability is $[1 - \Sigma'(E)]^{-1}$. It is at most one, as required. Eliminating a normalized eigenvector therefore changes its normalization metric as well as its spectral equation. Solving only a nonlinear energy equation does not provide that probability weight.

6.3 Exact elimination in the time domain

Theorem 6.3 (Causal memory with the initial hidden forcing)**.** *The finite system* $i\dot{u} = Au + Bv$, $i\dot{v} = Dv + B^*u$, *with initial data* (u_0, v_0), *is equivalent to*

$$\begin{aligned} i\dot{u}(t) &= Au(t) + Be^{-itD}v_0 - i\int_0^t K(t-s)u(s)\,ds, \\ K(\tau) &= Be^{-i\tau D}B^*, \qquad \tau \geq 0, \end{aligned} \tag{52}$$

together with reconstruction

$$v(t) = e^{-itD}v_0 - i\int_0^t e^{-i(t-s)D}B^*u(s)\,ds. \tag{53}$$

For $\operatorname{Im} z > 0$, *the self-energy is its exact one-sided transform:*

$$\Sigma(z) = -i\int_0^\infty e^{iz\tau}K(\tau)\,d\tau. \tag{54}$$

Proof. Differentiate $e^{itD}v(t)$ and integrate the hidden equation. This gives (53), whose substitution in the visible equation gives (52). Conversely the reconstruction differentiates to the hidden equation and supplies the prescribed initial value, so it recovers the full evolution. Since D is self-adjoint, $\|K(\tau)\| \leq \|B\|^2$. The factor $e^{-\operatorname{Im} z\,\tau}$ makes the transform absolutely convergent. On a spectral eigenvalue of D, $\int_0^\infty e^{i(z-\lambda)\tau}\,d\tau = i/(z-\lambda)$. The finite spectral sum then proves (54). □

The memory is causal because it uses only $s \leq t$. An initially empty hidden sector removes the forcing $Be^{-itD}v_0$, not the convolution. In finite dimension K is a finite sum of oscillatory terms. Its decay, a Markov approximation, or irreversible damping requires additional limits or assumptions. The closed full dynamics preserves $\|u(t)\|^2 + \|v(t)\|^2$, while the visible norm alone can decrease and later return. Energy-dependent Schur reduction is consequently different from replacing H by a fixed smaller self-adjoint matrix.

7 Rooted trees and radial Jacobi response

Let $T_{p,h}$ be a rooted p-ary tree of depth $h \geq 1$. Vertices at depth r have onsite energy $\epsilon_r \in \mathbb{R}$. Every parent-to-child edge from level r to $r+1$ has real coupling $\gamma_r/\sqrt{p}$, where $\gamma_r \in \mathbb{R}$. The Hamiltonian is the corresponding real symmetric adjacency matrix with onsite diagonal. The level symmetry and this branch normalization are structural assumptions of the model.

7.1 The radial reducing subspace

Define normalized level states

$$|r\rangle = p^{-r/2}\sum_{|v|=r}|v\rangle, \qquad r = 0, \ldots, h.$$

They are orthonormal because different levels are disjoint and level r contains p^r vertices. A direct action calculation gives

$$H|r\rangle = \gamma_{r-1}|r-1\rangle + \epsilon_r|r\rangle + \gamma_r|r+1\rangle, \tag{55}$$

with omitted terms at the endpoints. For example, the coefficient at each child at level $r+1$ is $p^{-r/2}\gamma_r/\sqrt{p} = \gamma_r p^{-(r+1)/2}$. At each parent at level $r-1$, the p child contributions instead give $p\,p^{-r/2}\gamma_{r-1}/\sqrt{p} = \gamma_{r-1}p^{-(r-1)/2}$.

Let $R : \mathbb{C}^{h+1} \to \ell^2(T_{p,h})$ send e_r to $|r\rangle$. Equation (55) proves $HR = RJ_h$, where

$$J_h = \begin{pmatrix} \epsilon_0 & \gamma_0 & & & \\ \gamma_0 & \epsilon_1 & \gamma_1 & & \\ & \gamma_1 & \epsilon_2 & \ddots & \\ & & \ddots & \ddots & \gamma_{h-1} \\ & & & \gamma_{h-1} & \epsilon_h \end{pmatrix}. \tag{56}$$

Because H is self-adjoint, the orthogonal complement of $\operatorname{Ran} R$ is also invariant: for $w \perp \operatorname{Ran} R$, $\langle Re_r, Hw\rangle = \langle HRe_r, w\rangle = 0$. Hence this is a reducing decomposition, rather than only a compressed approximation. The root belongs to the radial block.

7.2 Angular chains and their multiplicities

There is a more precise form of the complementary decomposition. For each vertex u at depth $a < h$, choose an orthonormal basis of vectors $c = (c_0, \dots, c_{p-1}) \in \mathbb{C}^p$ with $\sum_d c_d = 0$. For $s = 0, \dots, h-a-1$, set

$$|u, c; s\rangle = p^{-s/2} \sum_{d=0}^{p-1} c_d \sum_{w \in \{0,\dots,p-1\}^s} |udw\rangle. \tag{57}$$

For fixed u, c, these are unit vectors on successive levels. They obey the Jacobi recurrence with onsite energies $\epsilon_{a+1}, \dots, \epsilon_h$ and couplings $\gamma_{a+1}, \dots, \gamma_{h-1}$. At $s = 0$ the upward amplitude into u is proportional to $\sum_d c_d$, so it vanishes. At $s > 0$ the same branch-counting calculation as in (55) produces the neighboring chain vectors.

Proposition 7.1 (Complete angular-chain decomposition)**.** *The radial chain and all the chains* (57) *form an orthogonal decomposition of $\ell^2(T_{p,h})$. A chain starting at depth r has length $h-r+1$ and appears with multiplicity $(p-1)p^{r-1}$, for $1 \le r \le h$.*

Proof. Within one chain, different s have disjoint level supports. Chains based at the same u and different contrast vectors have inner product $\langle c, c'\rangle = 0$ at a common level. If two vertices are in disjoint subtrees, their chains have disjoint support. If one vertex lies below another, the higher chain is constant, up to its fixed ancestral coefficient, on all children below the lower vertex at a common level. The lower chain has zero sum over its own child contrasts, so their inner product is zero. The same zero-sum calculation makes every angular vector orthogonal to the uniform radial level vector. At a fixed depth b, the radial vector contributes one dimension. The angular vectors present there contribute $\sum_{r=1}^{b}(p-1)p^{r-1} = p^b - 1$ dimensions. They are orthogonal and therefore form a basis of that p^b-dimensional level space. Taking all levels proves completeness. There are p^{r-1} choices of the vertex u at depth $r-1$ and $p-1$ contrasts at each, giving the stated multiplicity. □

The rank $(p-1)p^{r-1}$ thus counts chains born at level r or new Haar directions on a level. A single level's detail space is usually transported to the next level; it is the entire angular chain that is invariant under the tree Hamiltonian. The full dark dimension is

$$\sum_{r=1}^{h}(p-1)p^{r-1}(h-r+1)=\sum_{b=0}^{h}p^b-(h+1).$$

Every one of these states is invisible to a root-localized symmetric source because it is orthogonal to the radial reducing block.

7.3 Continued fractions and their meromorphic reading

For a tail $J_{r:h}$, put $g_r(z)=\langle r,(z-J_{r:h})^{-1}r\rangle$. On a common regular domain, scalar Schur elimination gives

$$\begin{aligned} g_h(z)&=(z-\epsilon_h)^{-1},\\ g_r(z)&=\frac{1}{z-\epsilon_r-\gamma_r^2 g_{r+1}(z)}, \qquad r<h. \end{aligned} \tag{58}$$

Defining $\Sigma_h=0$ and $\Sigma_r=\gamma_r^2/(z-\epsilon_{r+1}-\Sigma_{r+1})$ gives

$$\langle \text{root},(z-H)^{-1}\text{root}\rangle=g_0(z)=\frac{1}{z-\epsilon_0-\Sigma_0(z)}. \tag{59}$$

The equality to the tree resolvent follows from the exact reducing decomposition. Equations (58) are rational identities. Nested denominators can become singular at a real energy where a canceled final rational expression is finite. One should use the resolvent or the canceled rational function at such a point, rather than perform division by zero in the nested chart.

The Jacobi coefficients in (56) do not depend on p. Thus the root Green function and radial propagator are prime independent when the same ϵ_r,γ_r are used. The complete tree spectrum retains the p-dependent angular multiplicities. If the per-edge coupling is instead a fixed g_r, the radial coupling is $\sqrt{p}\,g_r$ and its response depends on p. Prime independence is therefore a consequence of the declared normalization, not a universality statement about arbitrary trees.

7.4 A visible two-component block

As a conditional finite model, take $A=H_D(k)=k\sigma_z+m_0\sigma_x$ with $k,m_0\in\mathbb{R}$ and $D=J_{1:h}\otimes I_2$. Couple the visible spinor identically to the first site of the tail by $B=\gamma_0(\langle 1|\otimes I_2)$. Its self-energy is scalar in spin space:

$$\Sigma_D(z)=\sigma(z)I_2, \qquad \sigma(z)=\gamma_0^2\langle 1,(z-J_{1:h})^{-1}1\rangle.$$

The determinant formula from Theorem 6.1 is

$$\det(z-H_{\text{full}})=\det(z-D)\big((z-\sigma(z))^2-k^2-m_0^2\big). \tag{60}$$

Consequently any eigenvalue in the regular chart obeys

$$(E-\sigma(E))^2=k^2+m_0^2. \tag{61}$$

The formula is a dressed dispersion relation of this specific model. Its mass, couplings, and hidden sector have been supplied. No physical interaction or measured dispersion follows solely from the tree split.

For a nondegenerate visible spin direction with normalized low-space direction u_*, the visible weight of its normalized full eigenstate is $[1-\sigma'(E)]^{-1}$. If $k = m_0 = 0$, the spin eigenspace is degenerate, but the same normalization relation holds for each chosen unit vector in that eigenspace.

If every internal tail hopping is nonzero, the first tail site is cyclic. One can see this without an abstract cyclicity theorem: a tail eigenvector with zero first component has zero second component by the first row, and successive rows force all later components to vanish. Therefore every tail eigenvector has nonzero first component. At an eigenvalue λ of that tail, solvability of $(\lambda - D)v = B^* u$ requires its right side to be orthogonal to the tail eigenspace. With $\gamma_0 \neq 0$ and identical spin coupling this forces $u = 0$. The visible equation would then require $Bv = 0$, which forces $v = 0$ on the tail eigenspace. Thus in this irreducible case the full coupled spectrum avoids the tail eigenvalues. With a vanishing hopping, coincident visible and hidden energies can occur and the regular-chart qualification in (61) is essential.

The scalar causal kernel is $K(\tau) = \gamma_0^2 \langle 1, e^{-i\tau J_{1:h}} 1\rangle$, multiplying I_2. It uses the same radial coefficients as the rational self-energy. Its exact causal relation is (54).

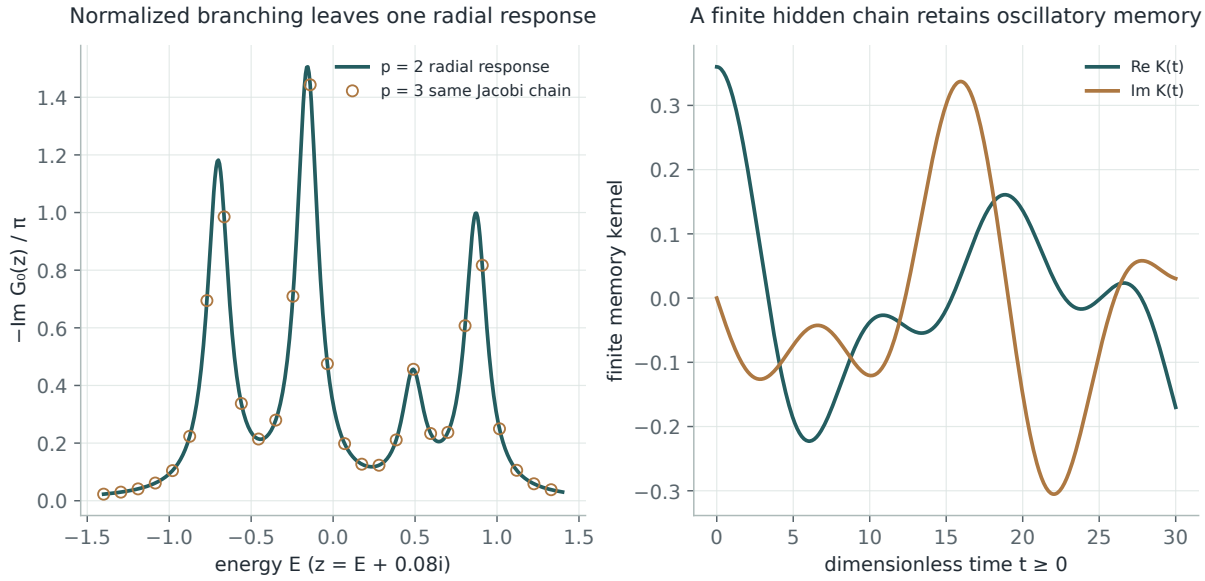


Figure 4. A branch-normalized finite tree has the same root response as its radial Jacobi chain. The complex-energy resolvent avoids real poles. The oscillatory memory is the return amplitude of a finite hidden chain.

8 Chebyshev decimation and resonant channels

We now eliminate intermediate residue sites in (44), assuming $\kappa \neq 0$. The retained sites are $C = \{0, p, \dots, (M-1)p\}$ and the eliminated sites comprise the $p-1$ interior sites between adjacent retained sites. Identify $\ell^2(C)$ with $\mathbb{C}^M$ and let $T|j\rangle = |j+1 \pmod M\rangle$. Set

$$x = \frac{z-\mu}{\kappa}, \qquad R_n(x) = 2T_n(x/2), \tag{62}$$

where T_n is the first-kind Chebyshev polynomial. The second-kind polynomials U_n satisfy $U_{-1} = 0$, $U_0 = 1$ and $U_{n+1}(t) = 2tU_n(t) - U_{n-1}(t)$.

8.1 The eliminated path fixes the Schur prefactor

The stationary equation on an interior site is $\psi_{n+1} = x\psi_n - \psi_{n-1}$. The eliminated block of $z - H$ in one cell is $\kappa K_{p-1}(x)$, where $K_n(x)$ has diagonal x and neighboring off-diagonal entries -1.

Lemma 8.1 (Dirichlet path determinant and endpoint entries)**.** *For $n \geq 1$, $\det K_n(x) = U_n(x/2)$. If this determinant is nonzero,*

$$(K_n(x)^{-1})_{11} = (K_n(x)^{-1})_{nn} = \frac{U_{n-1}(x/2)}{U_n(x/2)}, \qquad (K_n(x)^{-1})_{1n} = \frac{1}{U_n(x/2)}.$$

For $n = 1$ these entries all denote the same value $1/x$.

Proof. Expansion along the last row gives the determinant recurrence $d_n = xd_{n-1} - d_{n-2}$, with $d_0 = 1$ and $d_1 = x$. This is exactly the recurrence and initial data for $U_n(x/2)$. Deleting the first or last row and column leaves K_{n-1}, proving the diagonal inverse entries by cofactors. For the endpoint entry, the cofactor consists of the product of $n-1$ entries -1 along the unique path from the first endpoint to the last. Its cofactor sign is $(-1)^{n+1}$, so the product has total sign $(-1)^{n+1}(-1)^{n-1} = 1$. Division by d_n gives the stated value. □

Theorem 8.2 (Exact free-cycle decimation)**.** *For $U_{p-1}(x/2) \neq 0$, the Schur operator on the retained sublattice is*

$$\mathcal{S}_p(z) = \frac{\kappa}{U_{p-1}(x/2)}\big[R_p(x)I_M - (T + T^*)\big]. \tag{63}$$

Equivalently, $\mathcal{S}_p(z) = zI - H_{\text{eff}}^{(p)}(z)$, where

$$\begin{aligned} H_{\text{eff}}^{(p)}(z) &= \mu_p(z)I + \kappa_p(z)(T + T^*), \\ \kappa_p(z) &= \frac{\kappa}{U_{p-1}(x/2)}, \qquad \mu_p(z) = z - \frac{\kappa R_p(x)}{U_{p-1}(x/2)}. \end{aligned} \tag{64}$$

Proof. The eliminated sites split into M disjoint Dirichlet paths of length $p-1$. Hence their resolvent exists exactly when $U_{p-1}(x/2) \neq 0$. A retained site touches one endpoint on each of its two neighboring paths. Each path contributes the onsite self-energy $\kappa U_{p-2}/U_{p-1}$ by Lemma 8.1; the two contributions add. The path joining two adjacent retained sites contributes off-diagonal self-energy κ/U_{p-1}. Thus

$$\mathcal{S}_p(z) = \kappa\left(x - \frac{2U_{p-2}(x/2)}{U_{p-1}(x/2)}\right) I - \frac{\kappa}{U_{p-1}(x/2)}(T + T^*).$$

The identity $xU_{p-1}(x/2) - 2U_{p-2}(x/2) = 2T_p(x/2)$ follows from the three-term recurrence and $U_p - U_{p-2} = 2T_p$. For the latter identity, substituting $x/2 = \cos\theta$ gives $[\sin((p+1)\theta) - \sin((p-1)\theta)]/\sin\theta = 2\cos(p\theta)$, and polynomial equality extends it to all x. This proves (63) including its prefactor. Rearranging its diagonal gives (64). □

The normalized effective energy is $(z-\mu_p(z))/\kappa_p(z) = R_p(x)$. The parameter z remains fixed through the elimination. Since μ_p and κ_p depend on it, this is an energy-dependent spectral closure, not an energy-independent Hamiltonian flow. For $p = 2$, $R_2(x) = x^2 - 2$ and $U_1(x/2) = x$, giving

$$H^{(2)}_{\text{eff}}(z) = \mu I + \frac{\kappa^2}{z-\mu}(2I + T + T^*), \qquad z \neq \mu. \tag{65}$$

This decimation retains every other site. It does not retain the average of each neighboring pair.

8.2 The poles have an exact Dirichlet interpretation

The sine formula $U_{p-1}(\cos\theta) = \sin(p\theta)/\sin\theta$ shows that the excluded normalized energies are

$$x_j = 2\cos(j\pi/p), \qquad z_j = \mu + 2\kappa\cos(j\pi/p), \quad 1 \le j < p. \tag{66}$$

These are the energies of one eliminated path. Indeed its jth eigenvector has interior components $\sin(j\pi r/p)$, $r = 1, \dots, p-1$, and the nearest-neighbor recurrence gives precisely z_j. Both boundary values at $r = 0, p$ vanish. The eigenvectors are orthogonal and there are $p-1$ of them, so these exhaust the path spectrum.

The polynomial map R_p is defined at every x_j, while the Schur coefficients in (64) are not. Cancelling denominators in a determinant can cross this chart boundary; dividing by them in an eigenvector reconstruction cannot. A correct description there keeps the resonant cell modes or solves the full finite system. The fixed poles of the eliminated path are not all automatically eigenvalues of the full cyclic Hamiltonian.

Let Q project onto the eliminated sites. On the regular chart,

$$\det(z - H) = \det(z - QHQ)\det \mathcal{S}_p(z), \qquad \det(z - QHQ) = \left[\kappa^{p-1}U_{p-1}(x/2)\right]^M.$$

Their product is a polynomial identity after cancellation. It records both the hidden path factor and the coarse Schur factor, including their cancellations at the resonant energies.

8.3 Transfer coordinates and many-to-one branch maps

Define $A(x) = \begin{pmatrix} x & -1 \\ 1 & 0 \end{pmatrix}$. Its determinant is one and induction gives

$$A(x)^n = \begin{pmatrix} U_n(x/2) & -U_{n-1}(x/2) \\ U_{n-1}(x/2) & -U_{n-2}(x/2) \end{pmatrix}, \quad n \ge 1. \tag{67}$$

For $n = 1$ this uses $U_{-1} = 0$. Multiplication by $A(x)$ advances the first column through the Chebyshev recurrence and places the negative previous first column in the second. The trace is $R_n(x)$. Sampling the fine recurrence every p steps therefore uses the unimodular transfer matrix $A(x)^p$. Cayley–Hamilton gives the coarse recurrence $\chi_{j+1} - R_p(x)\chi_j + \chi_{j-1} = 0$ for $\chi_j = \psi_{pj}$. It explains the normalized spectral polynomial. The path-inverse proof above additionally supplies the prefactor needed for the resolvent.

The folded branches from Proposition 5.2, in normalized coordinates, satisfy

$$x_{p,q}(\theta) = 2\cos\left(\frac{\theta + 2\pi q}{p}\right), \qquad R_p(x_{p,q}(\theta)) = 2\cos\theta. \tag{68}$$

The second equality is just $T_p(\cos\alpha) = \cos(p\alpha)$. For a coarse normalized energy in the open interval $(-2, 2)$, choose its angle in $(0, \pi)$. These formulas supply the p distinct real roots of $R_p(x) = 2\cos\theta$. At the endpoints, some roots coincide and algebraic multiplicities must be retained. Globally the branches exchange labels when the angle crosses a cell boundary. There is no single-valued inverse of R_p on $[-2, 2]$.

For the fine cycle, $x_k = 2\cos(2\pi k/(pM))$ maps to $2\cos(2\pi k/M)$. Each coarse Fourier label has exactly p fine preimages. If two labels already have the same numerical energy, their multiplicities add as usual. Band folding retains the branch label and its vector, whereas normalized decimation identifies all of these labels at the same coarse energy. Their spectral relation does not make elimination an invertible operator change of basis.

9 Composition, transfer traces, and the free spectral measure

The spectral maps compose more simply than the energy-dependent operators that produce them. Keeping the normalized polynomial map, the rational Schur chart, and the transfer matrix distinct prevents an exact free-chain identity from being used outside its model.

9.1 Polynomial composition and iterated regular elimination

Theorem 9.1 (Chebyshev composition and Schur composition)**.** *For all positive integers m, n, the polynomial identities*

$$\begin{aligned} R_n(R_m(x)) &= R_{mn}(x) = R_m(R_n(x)), \\ U_{mn-1}(t) &= U_{n-1}(T_m(t))\, U_{m-1}(t) \end{aligned} \tag{69}$$

hold over $\mathbb{C}$. In a free cycle with $N = mnL$, $L \geq 3$, successive m- then n-site decimation at a fixed z agrees with direct mn-site decimation whenever its intermediate Schur inverses exist. Its final coefficients are

$$\kappa_{mn}(z) = \frac{\kappa}{U_{mn-1}(x/2)}, \qquad \mu_{mn}(z) = z - \frac{\kappa R_{mn}(x)}{U_{mn-1}(x/2)}. \tag{70}$$

Proof. For $x = 2\cos\theta$, first-kind composition gives $R_n(R_m(x)) = 2\cos(mn\theta) = R_{mn}(x)$. It holds on an interval and is therefore an identity of polynomials, including complex x. For $t = \cos\theta$ away from zeros of the sine denominators,

$$U_{n-1}(T_m(t))U_{m-1}(t) = \frac{\sin(mn\theta)}{\sin(m\theta)}\frac{\sin(m\theta)}{\sin\theta} = U_{mn-1}(t).$$

Polynomial equality extends it through those zeros and to all t. After the first regular elimination, the normalized variable is $x_m = (z - \mu_m)/\kappa_m = R_m(x)$. A second one gives

$$\kappa_{m,n} = \frac{\kappa_m}{U_{n-1}(x_m/2)} = \frac{\kappa}{U_{m-1}(x/2)U_{n-1}(T_m(x/2))} = \frac{\kappa}{U_{mn-1}(x/2)}.$$

Similarly $\mu_{m,n} = z - \kappa_{m,n}R_n(R_m(x))$, which gives (70). At each stage the matrices remain scalar nearest-neighbor circulants with the displayed complex coefficients, so their algebraic Schur formula continues to apply. These are rational identities on a common regular chart. They need not be interpreted as self-adjoint matrices at complex z. □

The product identity in (69) makes the exceptional denominators explicit. A nonzero final denominator implies both denominators along that two-stage chart are nonzero. At a zero of the final denominator the polynomial trace composition is still valid, but the divided Schur formulas require a new chart or retained resonant modes. Iterating one prime p through depth h, where the cycle has the required divisibility, replaces p by p^h in (70). Different prime orders give the same regular free-chain result. A nontranslation-invariant onsite potential, an interacting Hamiltonian, or a field-dependent constraint elimination needs its own closure analysis.

9.2 What remains true for a general unimodular transfer matrix

Let $M_0 \in SL_2(\mathbb{C})$ and $x = \operatorname{tr} M_0$. No diagonalizability assumption is needed. Cayley–Hamilton gives $M_0^2 - xM_0 + I = 0$, hence

$$M_0^n = U_{n-1}(x/2)M_0 - U_{n-2}(x/2)I, \qquad n \geq 1, \tag{71}$$

with $U_{-1} = 0$. For $n = 1$ it is the identity $M_0 = M_0$, and for $n = 2$ it is Cayley–Hamilton. Multiplying by M_0 and replacing M_0^2 by $xM_0 - I$ proves the induction step. Taking traces yields

$$\operatorname{tr}(M_0^n) = R_n(\operatorname{tr} M_0). \tag{72}$$

The coefficient identity is the same $xU_{n-1} - 2U_{n-2} = 2T_n$ used above.

If $w_{k+1} = M_0 w_k$, each residue subsequence satisfies

$$w_{nq+j} = M_0^j (M_0^n)^q w_0, \qquad 0 \leq j < n.$$

Thus ordinary matrix powering produces the normalized trace semigroup, including matrices with repeated eigenvalues at trace ± 2. The trace alone does not determine M_0, its eigenvectors, or a local Hamiltonian realization. In the free chain, $M_0 = A(x)$ has the required nearest-neighbor realization. For a general transfer problem only the matrix power and trace statements have been proved.

9.3 The arcsine measure and the Koopman operator

On $[-2, 2]$ define the probability measure

$$d\nu(x) = \frac{dx}{\pi\sqrt{4 - x^2}}. \tag{73}$$

It is the distribution of $2\cos\theta$ for a uniform angle on the circle, or on $[0, \pi]$. It is also the limiting normalized spectral measure of the free cycles: for a continuous test function f,

$$\frac{1}{N}\sum_{k=0}^{N-1} f(2\cos(2\pi k/N)) \longrightarrow \frac{1}{2\pi}\int_0^{2\pi} f(2\cos\theta)\, d\theta.$$

This is the Riemann-sum limit of a continuous periodic function.

Proposition 9.2 (Invariant free measure and exact mode routing). *For every integer $p \geq 2$, $(R_p)_*\nu = \nu$. The map $\mathcal{K}_p f = f \circ R_p$ is an isometry on $L^2(\nu)$. With*

$$\phi_0 = 1, \qquad \phi_n(x) = \sqrt{2}\, T_n(x/2), \quad n \geq 1,$$

its action and adjoint on an orthonormal basis are

$$\mathcal{K}_p \phi_n = \phi_{pn}, \qquad \mathcal{K}_p^* \phi_m = \begin{cases} \phi_{m/p}, & p \mid m, \\ 0, & p \nmid m. \end{cases} \tag{74}$$

In the first case $m = 0$ is included, so $\mathcal{K}_p^* \phi_0 = \phi_0$.

Proof. For a bounded measurable f, use the full circle parameterization:

$$\begin{aligned} \int f(R_p(x))\, d\nu(x) &= \frac{1}{2\pi} \int_0^{2\pi} f(2\cos(p\theta))\, d\theta \\ &= \frac{1}{2\pi p} \int_0^{2\pi p} f(2\cos u)\, du = \int f(x)\, d\nu(x). \end{aligned}$$

The last equality uses periodicity. Applying this equality to $|f|^2$ proves the isometry, first on bounded functions and then by truncation on all of $L^2(\nu)$. Pulling back to $[0, \pi]$ makes the ϕ_n the normalized cosine functions, and elementary cosine orthogonality proves their orthonormality. Their span is dense: Chebyshev polynomials span ordinary polynomials; those approximate continuous functions uniformly, for example by Bernstein polynomials after rescaling to $[0, 1]$. The approximation follows from uniform continuity and the binomial variance bound, by separating $|k/n - t| \leq \delta$ from its complement and then letting $n \to \infty$, $\delta \to 0$. Continuous functions are dense for this finite measure, as one sees by approximating interval steps in the angle coordinate by continuous ramps. This supplies completeness. Finally, Chebyshev composition gives $\mathcal{K}_p \phi_n = \phi_{pn}$. For every $n, m \geq 0$, $\langle \phi_n, \mathcal{K}_p^* \phi_m \rangle = \langle \phi_{pn}, \phi_m \rangle = \delta_{pn,m}$. Completeness yields the adjoint formula. □

$\mathcal{K}_p$ is an isometry with proper range, rather than a unitary: for example ϕ_1 is orthogonal to its range. Its adjoint discards nondivisible modes. Invariance of ν is a statement about the normalized free spectral distribution under R_p. It is not a proof that a sequence of growing Hilbert spaces has converged to an interacting Hamiltonian or that its quantum state converges. The direct-limit operator construction has separate hypotheses in the final part.

The mean logarithmic expansion of this map is also exact:

$$\int_{-2}^{2} \log |R_p'(x)|\, d\nu(x) = \log p. \tag{75}$$

Indeed for $x = 2\cos\theta$, $R_p'(x) = p \sin(p\theta)/\sin\theta$. Logarithms at its finitely many zeros are integrable, since they have the local form $\log|\theta - \theta_0|$ plus a bounded term. Hence the integral is $\log p$ plus the difference of the integrals of $\log|\sin(p\theta)|$ and $\log\sin\theta$ over $[0, \pi]$. The substitution $u = p\theta$ and π-periodicity of $\log|\sin u|$ make these two integrals equal. This calculation does not assert the entropy of a different physical dynamics.

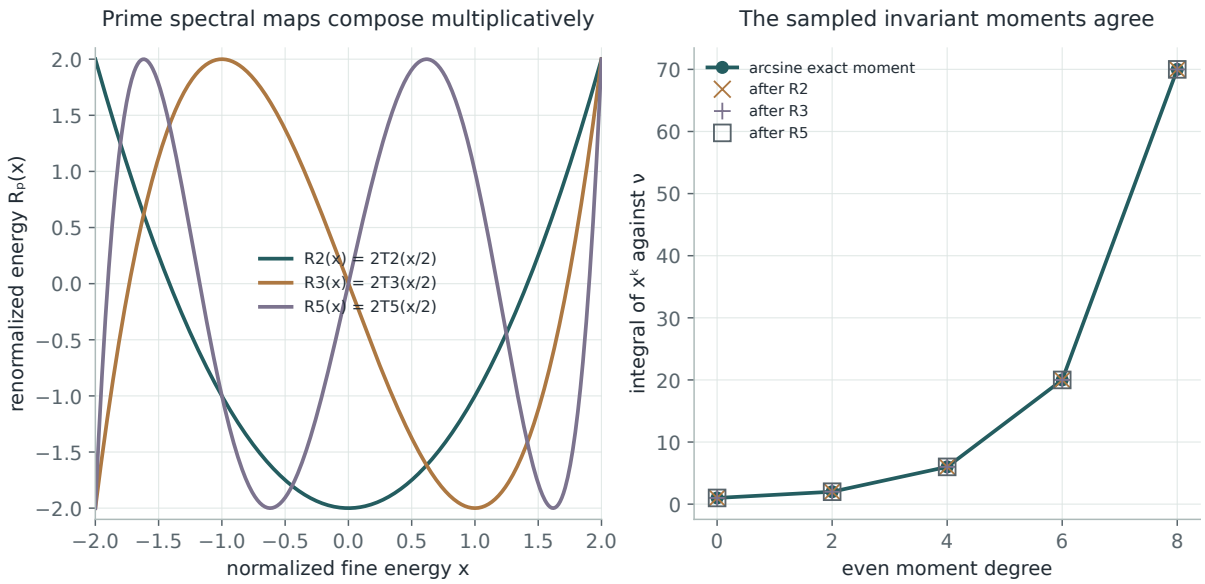


Figure 5. The normalized maps R_2, R_3, and R_5 preserve the same arcsine spectral measure. The accompanying bounded moment checks illustrate the exact measure identity, while each Schur realization retains its own excluded poles.

10 Analytic residue lanes and hypergeometric functions

Residue filtering also acts on coefficients of a convergent analytic series. Here the input is a function, not a quantum state, and the projector acts on its coefficient sequence. Evaluating one resulting lane at an operator does not generally turn that lane into a Hilbert space projector.

10.1 Character filtering and factorial separation

Fix an integer $L \geq 2$, write $\omega_L = e^{2\pi i/L}$, and let $F(z) = \sum_{n=0}^{\infty} a_n z^n$ be analytic for $|z| < \rho$. The whole disk is invariant under the rotations used below. Define

$$(\mathfrak{R}_{L,j}F)(z) = \frac{1}{L}\sum_{\ell=0}^{L-1} \omega_L^{-j\ell} F(\omega_L^{\ell} z), \quad 0 \leq j < L. \tag{76}$$

Absolute convergence for $|z| < \rho$ and the finite character sum give

$$(\mathfrak{R}_{L,j}F)(z) = \sum_{q=0}^{\infty} a_{Lq+j} z^{Lq+j}, \quad \mathfrak{R}_{L,j}\mathfrak{R}_{L,k} = \delta_{jk}\mathfrak{R}_{L,j}, \quad \sum_j \mathfrak{R}_{L,j} = I. \tag{77}$$

To prove the coefficient formula, the coefficient of $a_n z^n$ in (76) is $L^{-1}\sum_{\ell} \omega_L^{(n-j)\ell}$, equal to one exactly when $n \equiv j \pmod{L}$. The idempotent and reconstruction laws follow by filtering these coefficients again. For an entire F the formulas hold everywhere. For a function continued outside its disk, the domain and any chosen branches must also support all rotations; the local power-series identity does not remove such analytic restrictions.

Apply this filter to e^z and put

$$E_{L,j}(z) = \sum_{q=0}^{\infty} \frac{z^{Lq+j}}{(Lq+j)!}. \tag{78}$$

For $\alpha \in \mathbb{C}$, write $(\alpha)_q = \alpha(\alpha+1)\cdots(\alpha+q-1)$ and $(\alpha)_0 = 1$. Splitting the factorial into step-L factors gives

$$\begin{aligned} \frac{(Lq+j)!}{j!} &= \prod_{r=1}^{L} \prod_{k=0}^{q-1} (j+r+Lk) \\ &= L^{Lq} \prod_{r=1}^{L} \left(\frac{j+r}{L}\right)_q . \end{aligned} \tag{79}$$

Exactly one bottom parameter is one, at $r = L-j$. Define the multiset

$$\mathcal{A}_{L,j} = \left\{ \frac{j+r}{L} : 1 \le r \le L \right\} \setminus \{1\}.$$

Every member is positive, so none is a forbidden nonpositive integer in a hypergeometric denominator. Since $(1)_q = q!$,

$$E_{L,j}(z) = \frac{z^j}{j!}\, {}_0F_{L-1}\left(\begin{matrix} - \\ \mathcal{A}_{L,j} \end{matrix} ; \frac{z^L}{L^L} \right). \tag{80}$$

Here ${}_0F_s(;b_1,\dots,b_s;w)$ is defined by $\sum_{q\ge 0} w^q/[q!\prod_r (b_r)_q]$. The ratio of successive terms has magnitude $|w|/[(q+1)\prod_r |q+b_r|]$, which tends to zero. Its series is entire in w, so (80) is an everywhere convergent identity with the factorial series. For example $\mathcal{A}_{3,1} = \{2/3, 4/3\}$ and $E_{3,1}(z) = z\,{}_0F_2(;2/3,4/3;z^3/27)$.

10.2 The cyclic companion system

Termwise differentiation on compact subsets gives $E'_{L,j} = E_{L,j-1}$ for $j \ge 1$. For $j = 0$, the constant term vanishes under differentiation and reindexing gives $E'_{L,0} = E_{L,L-1}$. Let $\mathbf{E}_L = (E_{L,0},\dots,E_{L,L-1})^{\mathsf{T}}$ and define the cyclic matrix by $(C_L v)_0 = v_{L-1}$ and $(C_L v)_j = v_{j-1}$ for $j \ge 1$. Then

$$\mathbf{E}'_L = C_L \mathbf{E}_L, \qquad \mathbf{E}_L(0) = e_0, \qquad C_L^L = I, \qquad \mathbf{E}_L(z) = e^{zC_L} e_0. \tag{81}$$

One can verify the last equality directly: $C_L^n e_0 = e_{n \bmod L}$, so its jth exponential coordinate is exactly (78). Character filtering also gives

$$E_{L,j}(z) = \frac{1}{L} \sum_{\ell=0}^{L-1} \omega_L^{-j\ell} e^{\omega_L^{\ell} z}. \tag{82}$$

Thus factorial, hypergeometric, companion-state, and root-filter forms describe the same analytic function. Each form has a different numerical conditioning: a root sum can contain cancellations even when the defining series has positive terms.

For a general finite transfer matrix M_1, the residue-state identity $v_{Lq+j} = M_1^j (M_1^L)^q v_0$ follows by adding exponents. It needs no hypergeometric coefficients. A factorial lane provides a particular analytic coordinate system; it does not impose a unique microscopic Hamiltonian on every such transfer problem.

10.3 Bounded operators and controlled series errors

Theorem 10.1 (Entire lane calculus for a bounded operator)**.** *Let $A_0 \in \mathcal{B}(\mathcal{H})$ and $t \in \mathbb{C}$. The series $E_{L,j}(tA_0)$ converges absolutely in operator norm and satisfies*

$$\begin{aligned} E_{L,j}(tA_0) &= \frac{1}{L}\sum_{\ell=0}^{L-1} \omega_L^{-j\ell} e^{\omega_L^{\ell} tA_0}, \\ &= \frac{(tA_0)^j}{j!}\, {}_0F_{L-1}\left(\begin{matrix} - \\ \mathcal{A}_{L,j} \end{matrix}; \frac{(tA_0)^L}{L^L}\right), \\ \sum_{j=0}^{L-1} E_{L,j}(tA_0) &= e^{tA_0}. \end{aligned} \tag{83}$$

If all terms through degree Q are retained in a lane, the omitted operator norm is at most

$$e^a \frac{a^{Q+1}}{(Q+1)!}, \qquad a = |t|\,\|A_0\|\,. \tag{84}$$

Proof. The sum of norms of all exponential terms is at most e^a. A lane is a subseries of this sum, so it converges absolutely. The character calculation and factorial split can therefore be applied in operator norm, giving the first two identities. Partitioning the nonnegative degrees by their residues proves the third. All products are powers of the same operator, so there is no commuting-operator assumption. For the tail, bound a lane by the full tail. Since $(Q+1+k)! \geq (Q+1)!k!$,

$$\sum_{n=Q+1}^{\infty} \frac{a^n}{n!} \leq \frac{a^{Q+1}}{(Q+1)!} \sum_{k=0}^{\infty} \frac{a^k}{k!},$$

which proves (84). □

If $A_0^L = \alpha I$, the hypergeometric expression reduces every lane to a scalar function multiplying A_0^j. The propagator lies in the finite span of $I, A_0, \ldots, A_0^{L-1}$. A more general polynomial identity gives the corresponding finite algebra of powers. These statements follow from the supplied algebraic relation; they do not assert one for an arbitrary Hamiltonian.

The coefficient maps $\mathfrak{R}_{L,j}$ in (77) are idempotents on analytic functions. The evaluated operators $E_{L,j}(tA_0)$ generally are not. For $L = 2$, $E_{2,0}(tA_0) = \cosh(tA_0)$ and $E_{2,1}(tA_0) = \sinh(tA_0)$, which are neither an orthogonal projector family nor a pair of independent unitary evolutions. The sum reconstructs an evolution; an individual lane is an analytic contribution to it.

10.4 A squared-factorial lane for Bessel functions

The same separation gives the entire function

$$B(y) = \sum_{n=0}^{\infty} \frac{(-1)^n y^n}{(n!)^2} = {}_0F_1(;1;-y).$$

For $y \geq 0$, this is $J_0(2\sqrt{y})$; the entire series itself does not need a square-root branch. Define $B_{L,j}(y) = \sum_{q\geq 0}(-1)^{Lq+j}y^{Lq+j}/[(Lq+j)!]^2$. Squaring (79) gives

$$B_{L,j}(y) = \frac{(-1)^j y^j}{(j!)^2}\,{}_0F_{2L-1}\left(\begin{matrix} - \\ 1, \mathcal{A}_{L,j}, \mathcal{A}_{L,j}\end{matrix}; \frac{(-1)^L y^L}{L^{2L}}\right). \tag{85}$$

One of the two $q!$ factors is the conventional hypergeometric factorial and the other supplies the explicit bottom parameter 1. All bottom parameters are positive and the ratio test again gives entire convergence. Character filtering reconstructs $B = \sum_j B_{L,j}$. With the Euler derivative $\Theta = y\,d/dy$, termwise differentiation also gives the cyclic system

$$\Theta^2 B_{L,j}(y) = -yB_{L,j-1 \bmod L}(y). \tag{86}$$

For a term of positive degree n, its new coefficient is $n^2(-1)^n/(n!)^2 = -(-1)^{n-1}/[(n-1)!]^2$. The degree-zero term disappears under Θ^2, which also verifies the wraparound case $j = 0$. The scalar Gowdy Bessel mode therefore admits these analytic coordinates. Filtering that already-defined solution does not add a gravitational field equation.

11 Spectral domains and number-residue packets

The bounded calculus has two useful continuations. An unbounded Hamiltonian admits spectral functions with their exact domains. An oscillator admits actual residue projectors on its number spectrum. These are distinct constructions, and only the latter is a family of orthogonal quantum packet projectors.

11.1 An unbounded propagator does not make every lane bounded

Let H_∞ be self-adjoint with spectral measure E. For real t define the scalar functions $f_j(\lambda) = E_{L,j}(-it\lambda)$. The operator $f_j(H_\infty)$ has domain

$$\mathcal{D}(f_j(H_\infty)) = \left\{\psi : \int_{\mathbb{R}} |f_j(\lambda)|^2\, d\langle\psi, E(\lambda)\psi\rangle < \infty\right\}. \tag{87}$$

The scalar identity $\sum_j f_j(\lambda) = e^{-it\lambda}$ gives

$$\sum_j f_j(H_\infty)\psi = e^{-itH_\infty}\psi, \qquad \psi \in \bigcap_j \mathcal{D}(f_j(H_\infty)). \tag{88}$$

The equality holds on the stated common domain because it is a finite sum of spectral functions there. It does not say that each lane acts on every vector.

For an explicit example, take $L = 4$. The scalar formula gives

$$f_0(\lambda) = \tfrac{1}{2}(\cosh(t\lambda) + \cos(t\lambda)), \quad f_2(\lambda) = \tfrac{1}{2}(\cos(t\lambda) - \cosh(t\lambda)).$$

If H_∞ is multiplication by λ on $L^2(\mathbb{R})$ and $t \neq 0$, these functions define unbounded operators. Their sum cancels the growing hyperbolic contribution. The full propagator e^{-itH_∞} is nevertheless unitary on all of $L^2(\mathbb{R})$. For $L = 2$, by contrast, the two Schrödinger lanes are $\cos(tH_\infty)$ and $-i\sin(tH_\infty)$, both bounded. A larger lane count must not inherit this special dyadic boundedness without proof.

A sufficient common-domain condition is

$$\int_{\mathbb{R}} e^{2|t||\lambda|}\, d\langle\psi, E(\lambda)\psi\rangle < \infty, \tag{89}$$

since $|f_j(\lambda)| \le e^{|t||\lambda|}$. Spectral cutoffs $P_R = E([-R, R])$ supply a dense set satisfying this condition. On Ran P_R, the restricted Hamiltonian is bounded, the norm-series identities apply, and

$$\sum_j f_j(H_\infty) P_R \psi = e^{-itH_\infty} P_R \psi \longrightarrow e^{-itH_\infty}\psi.$$

On the common domain of the individual lanes, each term also converges to $f_j(H_\infty)\psi$ by its spectral-domain integrability. For an arbitrary vector only convergence of the summed cutoff expression is asserted. The domain question for an interacting Hamiltonian therefore remains even when a bounded surrogate has excellent numerical residuals.

11.2 A short factorial tail and its adic unit window

Let p be prime, let $L = mp$ for a positive integer m, and take $0 \le j < p$. For every $q \ge 0$,

$$g_q = \frac{j!(Lq)!}{(Lq+j)!} = \binom{Lq+j}{j}^{-1} \in \mathbb{Z}_p^\times. \tag{90}$$

Indeed $\binom{Lq+j}{j} = \prod_{r=1}^{j} (Lq + r)/r$. Each numerator and denominator is a p-adic unit because $Lq \equiv 0 \pmod p$ and $1 \le r < p$. For $j = 0$ the empty product is one. The proof also covers $p = 2$ in its short window. Outside this window, a factor can cross a multiple of p and the same proof gives no unit guarantee. This rational coefficient statement supplies no claim about the p-adic convergence of the complex exponential series.

Analytic refinement of a factorial lane is independent of this unit condition. Splitting $q = pr + a$, $0 \le a < p$, gives the exact identity

$$E_{L,j}(z) = \sum_{a=0}^{p-1} E_{pL,\,j+aL}(z), \qquad 0 \le j < L. \tag{91}$$

Absolute convergence permits this partition of the summation index. The labels $j + aL$ all lie in $[0, pL)$, but need not lie in the parent's short unit window. Exact complex analytic refinement and the local p-adic unit assertion therefore have separate hypotheses.

11.3 Actual orthogonal packets on the oscillator number spectrum

On the oscillator Fock space, let $\widehat{N}|n\rangle = n|n\rangle$ and define the bounded unitary $V = e^{2\pi i \widehat{N}/L}$. Its spectral residues are

$$Q_j^{(L)} = \frac{1}{L} \sum_{\ell=0}^{L-1} \omega_L^{-j\ell} V^\ell = \sum_{n \equiv j \pmod L} |n\rangle\langle n|. \tag{92}$$

On every number basis vector, the same finite character sum proves the second equality. These $Q_j^{(L)}$ are bounded orthogonal projectors, $Q_j Q_k = \delta_{jk} Q_j$, and $\sum_j Q_j = I$. No exponential norm series for the unbounded number operator is needed; its spectral calculus defines V.

For a coherent state

$$|\alpha\rangle = e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle, \qquad \mu = |\alpha|^2 \geq 0,$$

the packet probability is the Poisson residue sum

$$\begin{aligned} p_{L,j}(\mu) &= \langle \alpha, Q_j^{(L)} \alpha \rangle = e^{-\mu} E_{L,j}(\mu) \\ &= \frac{1}{L} \sum_{\ell=0}^{L-1} \omega_L^{-j\ell} e^{\mu(\omega_L^\ell - 1)}. \end{aligned} \tag{93}$$

The factorial series proves nonnegativity, including the limiting values $p_{L,0}(0) = 1$ and $p_{L,j}(0) = 0$ for $j > 0$. Partition of the Poisson series proves $\sum_j p_{L,j} = 1$. Their equivalent hypergeometric form is $e^{-\mu}$ times (80) at $z = \mu$.

Proposition 11.1 (Approach to uniform residue probability)**.** *For fixed $L \geq 2$ and $\mu \geq 0$,*

$$\left| p_{L,j}(\mu) - \frac{1}{L} \right| \leq \frac{1}{L} \sum_{\ell=1}^{L-1} e^{-\mu(1-\cos(2\pi\ell/L))} \leq \frac{L-1}{L} e^{-\mu(1-\cos(2\pi/L))}. \tag{94}$$

Proof. In (93), the $\ell = 0$ term is $1/L$. The modulus of every other term is $e^{\mu(\operatorname{Re}\omega_L^\ell - 1)}/L$. The triangle inequality gives the first bound. Among the nonzero residues, the largest cosine is $\cos(2\pi/L)$, giving the second. □

The exponent is positive for each fixed L, but decreases as $2\pi^2/L^2$ when L grows. The conclusion is therefore uniform in the finite packet label j at fixed L; it is not a depth-uniform estimate at arbitrarily large residue modulus.

Finally the annihilation operator routes adjacent residues:

$$Q_r^{(L)} a Q_s^{(L)} = 0 \quad \text{unless } r \equiv s - 1 \pmod{L}, \qquad Q_r^{(L)} a^\dagger Q_s^{(L)} = 0 \quad \text{unless } r \equiv s + 1 \pmod{L}. \tag{95}$$

On a finite number-state combination this follows from $a|n\rangle = \sqrt{n}\,|n-1\rangle$ and $a^\dagger|n\rangle = \sqrt{n+1}\,|n+1\rangle$. Each residue projector preserves $\mathcal{D}(a)$ and $\mathcal{D}(a^\dagger)$, because it only removes terms from their weighted number norm. Truncation of a vector in either domain converges in the corresponding graph norm, so the identities extend to those domains. For the isolated oscillator Hamiltonian, which is a function of $\widehat{N}$, the packet projectors commute with the evolution and the weights stay constant. A driving term containing ladder operators can transfer weight according to the routing rule. Selecting such a drive is further model data.

These packet probabilities are directly interpretable within the ordinary oscillator model. A physical interpretation of adic digits, hidden tree levels, or a preferred hypergeometric residue split requires an independently specified coupling and observable. The exact analytic and spectral coordinates derived here remain useful without making that identification.

Compatible dynamics and infinite limits

12 One Hilbert space behind all cylinder depths

An isometric refinement map compares two finite registers. An infinite quantum dynamics requires more: a common completed Hilbert space, a specified operator domain, and a topology in which the cutoffs converge. These requirements can all be met for the cylinder tower. They are separate requirements for a walk on a growing finite tree.

Fix a prime p. The compact additive group $\mathbb{Z}_p$ has normalized Haar measure μ_p, with

$$B_{m,a} = a + p^m\mathbb{Z}_p, \qquad 0 \le a < p^m, \qquad \mu_p(B_{m,a}) = p^{-m}.$$

Write $\mathcal{H}_p = L^2(\mathbb{Z}_p, \mu_p; \mathbb{C})$. We use the complex inner product $\langle f, g\rangle = \int \overline{f} g \, d\mu_p$, linear in its second argument. The p-adic coordinate belongs to the measurable space. Amplitudes and the Born probabilities use ordinary complex Hilbert space.

12.1 Counting amplitudes and Haar densities

At depth m, the finite space $\mathcal{H}_{p,m} = \ell^2(\mathbb{Z}/p^m\mathbb{Z})$ has counting inner product and orthonormal basis e_a. Define

$$\mathcal{U}_{p,m} e_a = p^{m/2} \mathbf{1}_{B_{m,a}}, \qquad \mathcal{V}_m = \operatorname{span}\{\mathbf{1}_{B_{m,a}} : 0 \le a < p^m\}. \tag{96}$$

The cylinders are disjoint. Consequently

$$\langle \mathcal{U}_{p,m} e_a, \mathcal{U}_{p,m} e_b \rangle = p^m \mu_p(B_{m,a} \cap B_{m,b}) = \delta_{ab}.$$

The same indicators span $\mathcal{V}_m$, so $\mathcal{U}_{p,m}$ is unitary onto this subspace. For $c = \sum_a c_a e_a$, its image has constant value $p^{m/2} c_a$ on $B_{m,a}$. In particular,

$$\int_{B_{m,a}} |\mathcal{U}_{p,m} c|^2 \, d\mu_p = |c_a|^2. \tag{97}$$

There is no extra probability factor after this normalization.

The finite refinement map is

$$V_{p,m} e_a = \frac{1}{\sqrt{p}} \sum_{d=0}^{p-1} e_{a+dp^m}.$$

Since $B_{m,a} = \bigsqcup_d B_{m+1,a+dp^m}$, it satisfies

$$\mathcal{U}_{p,m+1} V_{p,m} = \mathcal{U}_{p,m}. \tag{98}$$

Indeed, the coefficient on each child indicator is $p^{-1/2} p^{(m+1)/2} = p^{m/2}$, precisely the coefficient on the parent. Thus the refinement changes the address resolution of an embedded state without changing the function in $\mathcal{H}_p$.

An equivalent finite convention uses the weighted norm

$$\|f\|^2_{m,\mathrm{Haar}} = p^{-m} \sum_a |f(a)|^2.$$

Here $f(a)$ is a density amplitude rather than a counting amplitude, and the embedding into $\mathcal{H}_p$ sends f to $\sum_a f(a)\mathbf{1}_{B_{m,a}}$. Its refinement repeats the parent value on all children, without $p^{-1/2}$. The coordinate change $f(a) = p^{m/2}c_a$ identifies the two conventions. The factors p^{-m} in the measure and $p^{m/2}$ in the counting embedding are a single normalization expressed in two places.

12.2 Cylinder averaging as an orthogonal projection

For $f \in \mathcal{H}_p$, define

$$(\mathsf{E}_m f)(x) = p^m \int_{x+p^m\mathbb{Z}_p} f(y)\, d\mu_p(y). \tag{99}$$

The integral is defined since L^2 is contained in L^1 on this probability space. It produces a function constant on each depth-m cylinder.

Proposition 12.1 (Nested cylinder projections)**.** *The operator E_m is the orthogonal projection onto $\mathcal{V}_m$. For $m, n \geq 0$,*

$$\mathsf{E}_m^* = \mathsf{E}_m, \qquad \mathsf{E}_m^2 = \mathsf{E}_m, \qquad \mathsf{E}_m\mathsf{E}_n = \mathsf{E}_{\min(m,n)}. \tag{100}$$

Moreover $\mathsf{E}_m f \to f$ in L^2 for every $f \in \mathcal{H}_p$.

Proof. For a cylinder-constant g, the integral of $f - \mathsf{E}_m f$ over each depth-m cylinder is zero. Multiplying by the constant value of $\overline{g}$ there and summing shows $\langle g, f - \mathsf{E}_m f\rangle = 0$. The image is in $\mathcal{V}_m$, and averaging a function in $\mathcal{V}_m$ returns the function. These statements characterize the orthogonal projection and establish the first two identities. The spaces are nested, because every cylinder partitions into its children. Orthogonal projections onto nested subspaces satisfy the third identity: the larger projection fixes the smaller range, and the reverse order follows by taking adjoints.

For completeness, the union of these ranges is dense. Cylinders form a countable clopen base for the topology on $\mathbb{Z}_p$ and generate its Borel sigma-algebra. If $f \in L^2$ is orthogonal to every cylinder indicator, the finite complex measure $\nu(A) = \int_A f\, d\mu_p$ vanishes on cylinders and their finite unions. The class of sets on which ν is zero contains the whole space, is closed under complements, and is closed under countable disjoint unions. The elementary monotone-class extension from the cylinder algebra therefore gives $\nu(A) = 0$ for every Borel set. Applying this to the positive and negative level sets of the real and imaginary parts of f yields $f = 0$ almost everywhere. The orthogonal complement of $\bigcup_m \mathcal{V}_m$ is thus zero.

Given $\varepsilon > 0$, choose $g \in \mathcal{V}_M$ with $\|f - g\|_2 < \varepsilon$. For $m \geq M$, the projection fixes g, so $\|f - \mathsf{E}_m f\|_2 = \|(I - \mathsf{E}_m)(f - g)\|_2 \leq \|f - g\|_2 < \varepsilon$. This proves strong convergence without exchanging an infinite integral with a limit. □

The completion of the isometric direct system $(\mathcal{H}_{p,m}, V_{p,m})$ is therefore $\mathcal{H}_p$: equation (98) gives a single isometric map on the algebraic direct limit, and its image is the dense union of the $\mathcal{V}_m$. For $m \leq M$, with $V_{p,m\to M} = V_{p,M-1}\cdots V_{p,m}$, one also has

$$\mathsf{E}_m\big|_{\mathcal{V}_M} = \mathcal{U}_{p,M} V_{p,m\to M} V_{p,m\to M}^* \mathcal{U}_{p,M}^*. \tag{101}$$

Both sides are the orthogonal projection in $\mathcal{V}_M$ onto $\mathcal{V}_m$. This is the exact relation between the finite matrices and the common infinite Hilbert space.

12.3 Orthogonal details and a complete Haar basis

Use $\mathsf{F}_0 = \mathsf{E}_0$ and $\mathsf{F}_m = \mathsf{E}_m - \mathsf{E}_{m-1}$ for $m \geq 1$. These are the infinite depth details called D_m in the archived chapter. They correspond to the finite $D_m^{(M)}$ after the embedding $\mathcal{U}_{p,M}$. Equation (100) gives

$$\mathsf{F}_m \mathsf{F}_n = \delta_{mn} \mathsf{F}_m, \qquad \sum_{m=0}^{M} \mathsf{F}_m = \mathsf{E}_M. \tag{102}$$

For example, if $1 \leq m < n$, the four products in $(\mathsf{E}_m - \mathsf{E}_{m-1})(\mathsf{E}_n - \mathsf{E}_{n-1})$ are $\mathsf{E}_m, \mathsf{E}_m, \mathsf{E}_{m-1}, \mathsf{E}_{m-1}$ and cancel. For $m = n$ the same expansion is $\mathsf{E}_m - \mathsf{E}_{m-1}$. The cases involving $m = 0$ use $\mathsf{E}_0 \mathsf{E}_n = \mathsf{E}_0$.

Consequently

$$\mathcal{H}_p = \bigoplus_{m=0}^{\infty} \operatorname{Ran} \mathsf{F}_m, \qquad \|f\|_2^2 = \sum_{m=0}^{\infty} \|\mathsf{F}_m f\|_2^2 . \tag{103}$$

The sum of the projections is I in the strong operator topology. The detail ranks are $d_0 = 1$ and $d_m = (p-1)p^{m-1}$ for $m \geq 1$, the differences of the dimensions $p^m - p^{m-1}$.

For an explicit basis, put $\omega_p = \exp(2\pi i/p)$ and define

$$\psi_{m,a,q} = p^{(m-1)/2} \sum_{d=0}^{p-1} \omega_p^{qd} \mathbf{1}_{B_{m,a+dp^{m-1}}}, \qquad \begin{matrix} m \geq 1, \quad 0 \leq a < p^{m-1}, \\ 1 \leq q < p. \end{matrix} \tag{104}$$

Each has zero average over its parent, since $\sum_d \omega_p^{qd} = 0$, and lies in $\operatorname{Ran} \mathsf{F}_m$. Its norm squared is $p^{m-1} p\, p^{-m} = 1$. Distinct parents have disjoint support, while two characters at the same parent have inner product $p^{-1} \sum_d \omega_p^{(q'-q)d} = \delta_{qq'}$. The number of these orthonormal functions is exactly d_m, so they span that level. Together with $\psi_0 = 1$, they form a complete orthonormal basis by (103).

12.4 A statewise approximation rate

For $s > 0$, define the Haar regularity norm

$$\|f\|_{\mathfrak{H}_p^s}^2 = \sum_{m \geq 0} p^{2sm} \|\mathsf{F}_m f\|_2^2 , \qquad \mathfrak{H}_p^s = \{ f \in \mathcal{H}_p : \|f\|_{\mathfrak{H}_p^s} < \infty \}.$$

No identification with another Sobolev convention is required here. The definition itself gives the quantitative cutoff estimate.

Proposition 12.2 (Exact tail and regularity bound)**.** *For $f \in \mathcal{H}_p$ and every $M \geq 0$,*

$$\|(I - \mathsf{E}_M) f\|_2^2 = \sum_{m > M} \|\mathsf{F}_m f\|_2^2 . \tag{105}$$

If $f \in \mathfrak{H}_p^s$, then

$$\|(I - \mathsf{E}_M) f\|_2 \leq p^{-s(M+1)} \|f\|_{\mathfrak{H}_p^s} . \tag{106}$$

The convergence $\mathsf{E}_M \to I$ is strong and is not convergence in operator norm.

Proof. The omitted details are orthogonal, so Parseval gives the equality. For $m > M$, $p^{2sm} \geq p^{2s(M+1)}$. Multiplying the omitted summands by this lower bound proves the squared estimate, and taking square roots gives (106). At every finite M there is a unit vector in $\operatorname{Ran} \mathsf{F}_{M+1}$, on which $I - \mathsf{E}_M$ acts as the identity. Thus $\|I - \mathsf{E}_M\| = 1$ for every M. □

The rate is a bound on states in the displayed regularity class. For an arbitrary L^2 state the tail still tends to zero, but no common geometric rate follows. A finite array already in $\mathcal{V}_L$ has zero tail once $M \geq L$, which is a different statement from testing an infinite tail.

13 Self-adjoint dynamics and what a cutoff converges to

13.1 The maximal domain for level energies

Let $(\epsilon_m)_{m\geq 0}$ be any sequence of real numbers. Define

$$\mathcal{D}(H_\epsilon) = \left\{ f \in \mathcal{H}_p : \sum_{m\geq 0} |\epsilon_m|^2 \left\| \mathsf{F}_m f \right\|_2^2 < \infty \right\},$$
$$H_\epsilon f = \sum_{m\geq 0} \epsilon_m \mathsf{F}_m f. \tag{107}$$

The series converges in L^2 exactly on this domain. It is not necessary to bound the energies uniformly, but it is necessary to state the domain.

Theorem 13.1 (Real level energies define a self-adjoint operator)**.** *The operator in* (107) *is densely defined, closed, and self-adjoint. The finite cylinder functions form an operator core. It is bounded if and only if* $\sup_m |\epsilon_m| < \infty$*, and its spectrum is*

$$\sigma(H_\epsilon) = \overline{\{\epsilon_m : m \geq 0\}}. \tag{108}$$

Proof. The orthogonal decomposition (103) identifies f with the square-summable sequence $f_m = \mathsf{F}_m f$. Finite sequences belong to the domain, so it is dense. The real multipliers imply symmetry. To identify the adjoint, suppose $g \in \mathcal{D}(H_\epsilon^*)$ and write $k = H_\epsilon^* g$. Testing the adjoint identity against an arbitrary vector supported on level m gives $k_m = \epsilon_m g_m$. Since k is square summable, $\sum_m |\epsilon_m|^2 \left\| g_m \right\|^2 < \infty$. Thus $g \in \mathcal{D}(H_\epsilon)$ and the adjoint agrees with H_ϵ. The reverse inclusion is symmetry. This proves self-adjointness and therefore closedness.

For $f \in \mathcal{D}(H_\epsilon)$, the cylinder truncation $\mathsf{E}_M f$ satisfies both $\left\| f - \mathsf{E}_M f \right\|_2 \to 0$ and $\left\| H_\epsilon (f - \mathsf{E}_M f) \right\|_2 \to 0$, because the two squared norms are the tails of convergent nonnegative series. This is precisely density in the graph norm and proves the core assertion. If the energies have supremum $C < \infty$, Parseval gives $\left\| H_\epsilon f \right\|_2 \leq C \left\| f \right\|_2$. Conversely, a unit vector on level m has image norm $|\epsilon_m|$, so boundedness forces a finite supremum.

If z is outside the closed energy set, the sequence $(\epsilon_m - z)^{-1}$ is bounded and defines an inverse for $H_\epsilon - z$. Its image lies in the domain because $\epsilon_m/(\epsilon_m - z) = 1 + z/(\epsilon_m - z)$ is bounded as well. This proves that such z is in the resolvent set. An energy ϵ_m is an eigenvalue, while a limit point z has unit level vectors u_j with $\left\| (H_\epsilon - z) u_j \right\|_2 \to 0$. A bounded inverse would contradict this limit. Hence the closed energy set is exactly the spectrum. □

This proof also excludes a hidden choice of boundary condition at infinite depth. On the finite-cylinder core, the real level operator has the unique self-adjoint closure just described.

13.2 Explicit spectral calculus and strong continuity

For bounded Borel $F : \mathbb{R} \to \mathbb{C}$, define

$$F(H_\epsilon)f = \sum_{m \geq 0} F(\epsilon_m)\mathsf{F}_m f. \tag{109}$$

The series converges by Parseval and $\|F(H_\epsilon)f\|_2 \leq \sup_m |F(\epsilon_m)|\,\|f\|_2$. The same level decomposition verifies sums, products, and adjoints of these multipliers directly. In particular,

$$U(t) = e^{-itH_\epsilon} = \sum_{m \geq 0} e^{-it\epsilon_m}\mathsf{F}_m, \tag{110}$$

$$(z - H_\epsilon)^{-1} = \sum_{m \geq 0} (z - \epsilon_m)^{-1}\mathsf{F}_m, \quad z \notin \sigma(H_\epsilon). \tag{111}$$

The projections in these sums converge strongly, not absolutely in operator norm in general.

The modulus-one multipliers give $U(t)^*U(t) = I$ and $U(t+s) = U(t)U(s)$. They also give strong continuity:

$$\|(U(t) - U(t_0))f\|_2^2 = \sum_m |e^{-it\epsilon_m} - e^{-it_0\epsilon_m}|^2\,\|\mathsf{F}_m f\|_2^2 \longrightarrow 0.$$

To justify this limit, first bound the omitted summands by $4\sum_{m>M}\|\mathsf{F}_m f\|_2^2$, then use ordinary continuity in the finite sum up to M. On $\mathcal{D}(H_\epsilon)$ the same argument, using $|(e^{-it\epsilon} - 1)/t| \leq |\epsilon|$, proves

$$\lim_{t \to 0} \frac{U(t)f - f}{t} = -iH_\epsilon f.$$

Thus the declared self-adjoint operator is the generator. A general L^2 state evolves unitarily even if this derivative does not exist.

13.3 Two distinct extensions of the finite dynamics

Let $H_\epsilon^{(M)}$ be the finite restriction of H_ϵ to $\mathcal{V}_M$. One useful approximation to $U(t)$ is its embedded finite propagator, extended by zero on the omitted subspace:

$$U_M^{\mathrm{emb}}(t) = \sum_{m=0}^{M} e^{-it\epsilon_m}\mathsf{F}_m = U(t)\mathsf{E}_M. \tag{112}$$

It is unitary on $\mathcal{V}_M$ and has kernel $\mathcal{V}_M^\perp$ as an operator on $\mathcal{H}_p$.

Proposition 13.2 (Exact state error for a compatible cutoff)**.** *For every $f \in \mathcal{H}_p$,*

$$\left\|U(t)f - U_M^{\mathrm{emb}}(t)f\right\|_2 = \left\|(I - \mathsf{E}_M)f\right\|_2 . \tag{113}$$

This equality is uniform for all real t. For $f \in \mathfrak{H}_p^s$, the right side is bounded by $p^{-s(M+1)} \left\|f\right\|_{\mathfrak{H}_p^s}$. For $f \in \mathcal{D}(H_\epsilon)$ there is also the exact graph-norm identity

$$\begin{aligned} &\left\|U(t)f - U_M^{\mathrm{emb}}(t)f\right\|_2^2 + \left\|H_\epsilon(U(t)f - U_M^{\mathrm{emb}}(t)f)\right\|_2^2 \\ &\quad = \sum_{m>M} (1 + |\epsilon_m|^2) \left\|\mathsf{F}_m f\right\|_2^2 . \end{aligned} \tag{114}$$

Proof. The difference is the omitted orthogonal sum $\sum_{m>M} e^{-it\epsilon_m}\mathsf{F}_m f$. Its multipliers have modulus one, so Parseval gives (113). Applying H_ϵ multiplies the omitted level by ϵ_m and gives the second identity. The domain hypothesis ensures that this second series is finite. The regularity estimate is Proposition 12.2. □

An alternative is to extend the finite Hamiltonian by the zero Hamiltonian on the omitted subspace:

$$\widetilde{H}_M = H_\epsilon \mathsf{E}_M = \sum_{m \le M} \epsilon_m \mathsf{F}_m .$$

This is a bounded self-adjoint operator on all of $\mathcal{H}_p$. Its propagator is a full-space unitary, but has a different tail:

$$e^{-it\widetilde{H}_M} = U_M^{\mathrm{emb}}(t) + (I - \mathsf{E}_M), \qquad \left\|(e^{-it\widetilde{H}_M} - U(t))f\right\|_2 \le 2 \left\|(I - \mathsf{E}_M)f\right\|_2 . \tag{115}$$

Likewise, for nonreal z,

$$\begin{aligned} (z - \widetilde{H}_M)^{-1} &= \mathsf{E}_M (z - H_\epsilon)^{-1} + z^{-1}(I - \mathsf{E}_M), \\ \left\|((z - \widetilde{H}_M)^{-1} - (z - H_\epsilon)^{-1})f\right\|_2 &\le (|z|^{-1} + |\operatorname{Im} z|^{-1}) \left\|(I - \mathsf{E}_M)f\right\|_2 . \end{aligned} \tag{116}$$

Thus $\widetilde{H}_M$ converges to H_ϵ in the strong resolvent sense. These statements specify operators on one common Hilbert space. They do not compare matrices of changing dimensions without an embedding.

The embedded propagators converge strongly, uniformly in time on each fixed state. Their operator norm error nevertheless satisfies

$$\left\|U(t) - U_M^{\mathrm{emb}}(t)\right\| = 1 \qquad \text{for every } M, t.$$

The fully unitary extensions also converge strongly with the bound in (115); norm convergence of these extensions needs a separate condition on the omitted phases. Neither assertion is a general norm-limit theorem for arbitrary quantum dynamics.

13.4 A graph-domain example with exact infinite tails

Example 13.3 (Convergent states and finite energy derivatives)**.** Choose one unit vector $u_m \in \operatorname{Ran} \mathsf{F}_m$ at every level, let $0 < q < 1$, and set

$$f_q = \sum_{m \ge 0} \sqrt{(1-q)q^m}\, u_m, \qquad \epsilon_m = 2^m .$$

Then $\|f_q\|_2 = 1$ and

$$\|(I - \mathsf{E}_M) f_q\|_2^2 = q^{M+1}. \tag{117}$$

However, $f_q \in \mathcal{D}(H_\epsilon)$ exactly when $4q < 1$:

$$\|H_\epsilon f_q\|_2^2 = (1-q) \sum_{m \geq 0} (4q)^m.$$

For $q < 1/4$, its graph tail is exactly

$$q^{M+1} + \frac{(1-q)(4q)^{M+1}}{1-4q}. \tag{118}$$

For $q = 1/8$, the state norm tail decays faster than the graph norm tail. For $q = 1/2$, the normalized state still has a convergent L^2 cylinder approximation and a well-defined unitary evolution, but is outside the Hamiltonian domain. All these conclusions concern genuinely infinite series, rather than a finite last level with its tail set to zero.

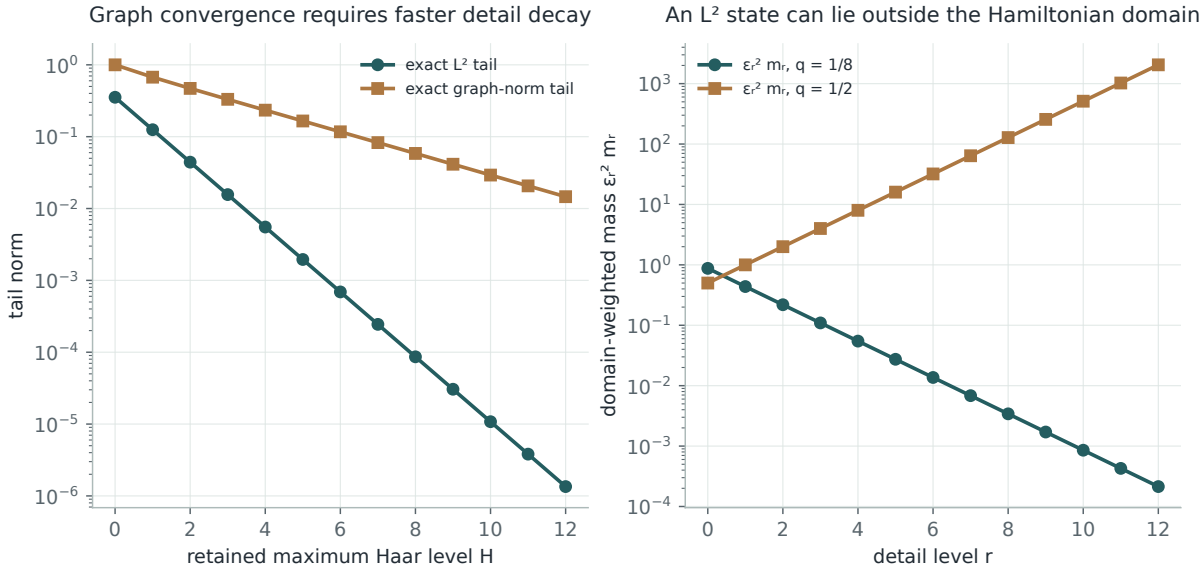


Figure 6. Exact infinite tails for the geometric level masses of Example 13.3. The operator domain controls a weighted tail in addition to the state norm tail. The parameters are declared in the plate and its reproducible Python source.

13.5 Compact resolvents and thermal weights

The omitted embedded resolvent has the exact norm

$$\|(I - \mathsf{E}_M)(z - H_\epsilon)^{-1}\| = \sup_{m > M} |z - \epsilon_m|^{-1}. \tag{119}$$

The upper bound follows by Parseval. For the lower bound, test on a unit vector at any omitted level and take the supremum. Since every level is nonzero and finite dimensional, this proves the following equivalence.

Proposition 13.4 (Exact compactness criterion). *The operator H_ϵ has compact resolvent if and only if $|\epsilon_m| \to \infty$. In that case the embedded finite resolvents converge in operator norm. This is an assertion about the embedded resolvents, rather than the zero-Hamiltonian extensions of* (116).

Proof. If the absolute energies tend to infinity, the right side of (119) tends to zero. The retained part of the resolvent has finite rank, so the resolvent is a norm limit of finite-rank operators and is compact. Conversely, if the absolute energies do not tend to infinity, there are infinitely many distinct levels m_j with $|\epsilon_{m_j}| \le C$ for some C. Unit vectors u_j on these levels are orthogonal. For a fixed nonreal z, their resolvent images are also orthogonal and have norm at least $(|z| + C)^{-1}$. They have no convergent subsequence. The resolvent is therefore not compact. □

When $|\epsilon_m| \to \infty$, the omitted part of the resolvent for $\widetilde{H}_M$ in (116) still carries the constant multiplier z^{-1}. It does not converge in norm to the omitted resolvent of H_ϵ. Strong resolvent convergence and norm convergence of an embedded finite-rank resolvent are distinct claims.

For $\beta > 0$, the multiplier $e^{-\beta H_\epsilon}$ is bounded precisely when the energies are bounded below. If there is no lower bound, the levelwise exponential is an unbounded operator with domain $\{f : \sum_m e^{-2\beta\epsilon_m} \|\mathsf{F}_m f\|_2^2 < \infty\}$. Thus a finite Euclidean intertwining identity does not by itself supply a bounded infinite heat semigroup. When bounded below, the positive operator has trace

$$\operatorname{Tr}(e^{-\beta H_\epsilon}) = e^{-\beta\epsilon_0} + \sum_{m\ge 1} (p-1)p^{m-1} e^{-\beta\epsilon_m}, \tag{120}$$

possibly infinite. This follows by summing its nonnegative diagonal entries in the complete Haar basis. It is trace class exactly when this series converges. For $\epsilon_m = cm$ with $c > 0$, the condition is $\beta c > \log p$, by the geometric-series test. For $\epsilon_m = p^{\alpha m}$ with $\alpha > 0$, it converges for every $\beta > 0$, since the exponential energy term eventually dominates the linear-in-m logarithm of the multiplicity. A thermodynamic partition function therefore needs an energy choice and a summability condition in addition to Hilbert-space reconstruction.

14 Compatible interactions across finitely many primes

14.1 Tensor embeddings and rectangular cutoffs

Let $\mathcal{P}$ be a finite prime set and let $\mathcal{K}_{\text{tor}}$ be a finite-dimensional complex space for the retained torsion and dyadic sign registers. Set

$$\mathcal{H}_{\mathcal{P}}^{\infty} = \mathcal{K}_{\text{tor}} \otimes \bigotimes_{p\in\mathcal{P}} L^2(\mathbb{Z}_p, \mu_p). \tag{121}$$

The corresponding finite register embeds by $I_{\mathcal{K}_{\text{tor}}} \otimes \bigotimes_p \mathcal{U}_{p,h_p}$. For a multi-depth $\mathbf{h} = (h_p)$, define

$$\mathsf{E}_{\mathbf{h}} = I_{\mathcal{K}_{\text{tor}}} \otimes \bigotimes_p \mathsf{E}_{p,h_p}.$$

Factorwise projection identities give

$$\mathsf{E}_{\mathbf{h}}\mathsf{E}_{\mathbf{k}} = \mathsf{E}_{\min(\mathbf{h},\mathbf{k})}, \qquad \mathsf{E}_{\mathbf{h}} f \to f \quad \text{as every } h_p \to \infty. \tag{122}$$

To justify the second statement, first take an elementary tensor. The difference of the tensor projections and the identity is a finite telescoping sum, one factor of which tends strongly to zero in each term,

while all other projection factors have norm at most one. Finite sums of elementary tensors are dense, and the projections are contractions. This extends convergence to every f.

Refinement in distinct coordinates is path independent because

$$(V_p \otimes I)(I \otimes V_q) = (I \otimes V_q)(V_p \otimes I) = V_p \otimes V_q.$$

The identity on $\mathcal{K}_{\text{tor}}$ leaves the separately retained sign channel untouched by the averaging. This is a finite set of primes. An infinite tensor product, a restricted adelic product, and a model including the Archimedean place each require further choices.

Define the joint detail projections

$$\mathsf{F}_{\mathbf{r}} = I_{\mathcal{K}_{\text{tor}}} \otimes \bigotimes_{p \in \mathcal{P}} \mathsf{F}_{p,r_p}, \qquad \mathbf{r} \in \mathbb{N}^{\mathcal{P}}. \tag{123}$$

They are pairwise orthogonal and have finite-dimensional ranges. Rectangular sums equal $\mathsf{E}_{\mathbf{h}}$, so their strong sum is the identity. This permits interactions between primes while preserving the exact depth filtration.

14.2 A complete block-domain theorem

The scalar weights are one particularly simple choice. The following editorial extension states the exact additional input for interacting blocks. It is useful because the torsion operators inside a joint packet need not commute with one another.

Theorem 14.1 (Self-adjoint packet blocks and compatible finite dynamics)**.** *For every* $\mathbf{r}$*, let* $A_{\mathbf{r}}$ *be a self-adjoint operator on the finite-dimensional space* $\mathcal{K}_{\mathbf{r}} = \operatorname{Ran} \mathsf{F}_{\mathbf{r}}$*. On* $\mathcal{H}_{\mathcal{P}}^{\infty} = \bigoplus_{\mathbf{r}} \mathcal{K}_{\mathbf{r}}$ *define*

$$\mathcal{D}(H_A) = \left\{ f : \sum_{\mathbf{r}} \|A_{\mathbf{r}} \mathsf{F}_{\mathbf{r}} f\|^2 < \infty \right\},$$
$$H_A f = \sum_{\mathbf{r}} A_{\mathbf{r}} \mathsf{F}_{\mathbf{r}} f. \tag{124}$$

Then H_A *is self-adjoint, and finite rectangular packet sums form a core. Every rectangular range reduces* H_A*, and its finite restriction* $H_A^{(\mathbf{h})}$ *satisfies exact embedding intertwining with every larger rectangular restriction. For bounded Borel* F*,*

$$F(H_A) f = \sum_{\mathbf{r}} F(A_{\mathbf{r}}) \mathsf{F}_{\mathbf{r}} f, \qquad F(H_A)\mathsf{E}_{\mathbf{h}} \to F(H_A) \quad \textit{strongly}. \tag{125}$$

In particular,

$$\left\| e^{-itH_A}(I - \mathsf{E}_{\mathbf{h}}) f \right\| = \left\| (I - \mathsf{E}_{\mathbf{h}}) f \right\|. \tag{126}$$

Proof. Finite rectangular sums belong to the domain and are dense. Symmetry holds by symmetry on every block. If g belongs to the adjoint domain and $k = H_A^* g$, testing against an arbitrary vector in $\mathcal{K}_{\mathbf{r}}$ gives $\mathsf{F}_{\mathbf{r}} k = A_{\mathbf{r}} \mathsf{F}_{\mathbf{r}} g$. Square summability of k is precisely the domain condition in (124). Thus $H_A^* = H_A$. Truncating the two square-summable sequences $f_{\mathbf{r}}$ and $A_{\mathbf{r}} f_{\mathbf{r}}$ proves the graph-core assertion.

The rectangular projection selects a collection of whole blocks, so it commutes with H_A on the domain and reduces the operator. The restriction to each retained block is literally the same matrix at every subsequent depth. This proves the finite intertwining, and finite-dimensional spectral calculus

on each block gives (125). Boundedness of F bounds its omitted image by $\|F\|_\infty \|(I - \mathsf{E}_{\mathbf{h}})f\|$, proving strong convergence. Finally, each block exponential is unitary, so orthogonality of the omitted blocks gives (126). □

For example, on a joint depth packet one may take

$$A_{\mathbf{r}} = \left(\sum_a \epsilon_{a,\mathbf{r}} T_a \right) \otimes I_{\text{phase},\mathbf{r}}, \qquad \epsilon_{a,\mathbf{r}} \in \mathbb{R}, \quad T_a = T_a^*.$$

The matrices T_a act on the sign and torsion factors. Their mutual commutation is unnecessary because their sum is self-adjoint on the finite block. The coefficients assigned to an existing packet are fixed under refinement. An extension acting wholly inside a new detail complement has no matrix element between that complement and the old range. A cross term between them would be a different model.

There is also an exact construction for any *specified* self-adjoint compatible finite family along a cofinal nested sequence. Suppose finite-dimensional spaces $\mathcal{V}_0 \subset \mathcal{V}_1 \subset \cdots$ have dense union, and self-adjoint K_M satisfy $K_{M+1} f = K_M f$ for $f \in \mathcal{V}_M$. Self-adjointness implies that $\mathcal{V}_M^\perp \cap \mathcal{V}_{M+1}$ is invariant: for g there and $f \in \mathcal{V}_M$, $\langle f, K_{M+1} g\rangle = \langle K_{M+1} f, g\rangle = 0$. The initial space $\mathcal{V}_0$ and the successive difference spaces therefore have stable self-adjoint finite blocks. Their orthogonal direct sum with the maximal square-summability domain is self-adjoint by the same adjoint proof above. This is an editorial compatible-family construction, not an assertion that arbitrary finite Hamiltonians are compatible. The identities $K_{M+1}|_{\mathcal{V}_M} = K_M$ are its substantive hypothesis.

14.3 An exact cross-prime entangling generator

Let Z_{sign} have eigenvalues $+1, -1$ on a dyadic sign qubit, and let $B = B^*$ act on the remaining factors. Assume there are orthonormal states $|b_0\rangle, |b_1\rangle$ with $B|b_0\rangle = 0$ and $B|b_1\rangle = |b_1\rangle$. One concrete choice is an odd-prime torsion involution times a fixed odd-prime detail projector, with the torsion state in its $+1$ eigenspace. For real $g \neq 0$, use

$$H_{\text{int}} = g Z_{\text{sign}} \otimes B, \qquad |\Psi(0)\rangle = |+_x\rangle \otimes \frac{|b_0\rangle + |b_1\rangle}{\sqrt{2}}.$$

Evolution under this explicit interaction alone gives

$$|\Psi(t)\rangle = \frac{1}{\sqrt{2}} \left(|+_x\rangle \otimes |b_0\rangle + e^{-igtZ_{\text{sign}}} |+_x\rangle \otimes |b_1\rangle \right). \tag{127}$$

Tracing out the orthogonal b states yields the equally weighted mixture of the two displayed sign states. Their overlap is $\cos(gt)$. A mixture of two unit vectors with equal weights and overlap magnitude c has eigenvalues $(1 + c)/2$ and $(1 - c)/2$: its trace is one and its determinant on their two-dimensional span is $(1 - c^2)/4$. Thus the sign entropy in bits is

$$S_{\text{sign}}(t) = h_2 \left(\frac{1 + |\cos(gt)|}{2} \right), \quad h_2(u) = -u \log_2 u - (1 - u) \log_2 (1 - u), \tag{128}$$

with $0 \log_2 0 = 0$. At $t_* = \pi/(2|g|)$ the two sign states are orthogonal and the entropy is one bit. A fixed level projector is compatible under refinement, so this interaction can entangle the prime factors while preserving the cylinder embedding. The tensor factorization alone creates no entanglement. It is this chosen generator and initial superposition that produce the result.

15 Controlled leakage when compatibility is broken

The preceding identities are exact because the coarse range reduces the Hamiltonian. A weighted coupling to its complement has a different, quantifiable behavior. This section concerns finite-dimensional self-adjoint operators; no unbounded perturbation theorem is implicit.

Let V be an isometry from a coarse space to a finite fine space, set $P = VV^*$ and $Q = I - P$, and let $H_0 = H_0^*$ satisfy $H_0V = VH_c$. Let $B = B^*$ be purely off diagonal:

$$PBP = 0, \qquad QBQ = 0.$$

For a real λ, define $H_\lambda = H_0 + \lambda B$. In the source model $\lambda = p^{-w}$ for an assigned weight $w \geq 1$.

Proposition 15.1 (Defect, leakage, and a second-order self-energy)**.** *Under these hypotheses,*

$$\|H_\lambda V - VH_c\| = |\lambda| \, \|BV\| , \tag{129}$$

$$\left\|Qe^{-itH_\lambda}P\right\| \leq \min\{1, |t\lambda| \, \|QBP\|\}. \tag{130}$$

For $z \notin \sigma(QH_\lambda Q|_{Q\mathcal{H}})$, the coarse self-energy is

$$\Sigma_P(z) = PH_\lambda Q(z - QH_\lambda Q|_{Q\mathcal{H}})^{-1} QH_\lambda P, \tag{131}$$

where the inverse acts only on $Q\mathcal{H}$, and

$$\|\Sigma_P(z)\| \leq \frac{|\lambda|^2 \, \|PBQ\|^2}{\operatorname{dist}(z, \sigma(QH_\lambda Q|_{Q\mathcal{H}}))}. \tag{132}$$

Proof. The defect is λBV, so the first identity is immediate. Self-adjointness of H_0 makes its invariant coarse range reducing. Hence $Qe^{-itH_0}P = 0$, and Duhamel's identity gives

$$Qe^{-itH_\lambda}P = -i\lambda \int_0^t Qe^{-i(t-s)H_\lambda} Be^{-isH_0}P \, ds.$$

The off-diagonal assumption and reduction by H_0 give $Be^{-isH_0}P = QBPe^{-isH_0}P$. The two propagators and the exterior projections are contractions. Integrating the norm bound over an interval of length $|t|$ gives the second expression in (130). The bound one follows directly from unitarity and the norms of the projections.

The coarse/detail blocks of H_0 vanish, so $PH_\lambda Q = \lambda PBQ$ and $QH_\lambda P = \lambda QBP$. The detail compression is self-adjoint. Diagonalizing it bounds the resolvent norm by the reciprocal spectral distance. Since $QBP = (PBQ)^*$, submultiplicativity gives (132). □

The small-time leading term is $-it\lambda QBP$. More precisely,

$$\left\|Qe^{-itH_\lambda}P + it\lambda QBP\right\| \leq \frac{|\lambda|t^2}{2}(\|H_\lambda\| + \|H_0\|) \, \|QBP\| . \tag{133}$$

To verify this, subtract the constant integrand QBP in Duhamel's formula and use $\left\|e^{-iuK} - I\right\| \leq |u| \, \|K\|$ for a finite self-adjoint K. The two deviations are bounded by $(|t-s| \, \|H_\lambda\| + |s| \, \|H_0\|) \, \|QBP\|$ on the oriented integration interval. Their integrals are both $t^2/2$. For $|\lambda| \leq 1$, this is a uniform $O(t^2|\lambda|)$

remainder. A generic $O(t^2\lambda^2)$ remainder would omit the first-order interaction with the background Hamiltonian.

For the stipulated assignment $\lambda = p^{-w}$, the operator defect and this leakage upper bound scale linearly with p^{-w}. The self-energy bound is quadratic in p^{-w} away from the detail spectrum. Under the purely off-diagonal hypothesis, $QH_\lambda Q = QH_0Q$, so the denominator does not depend on λ. Approaching a detail resonance removes that uniform estimate. These are exact finite quantum statements for a chosen coupling law. An integral valuation packet does not, by itself, force every physical interaction strength to equal p^{-w}.

16 A local walk on physical position and refinement depth

16.1 The finite graph and its unitary tick

Let $N \geq 3$ and $h \geq 0$. Let C_N be the physical cycle and let $T_{p,h}$ be the rooted finite p-ary tree of prefixes of length at most h. The product graph is

$$\mathcal{G} = C_N \square T_{p,h}.$$

A vertex is $v = (x, \mathbf{d})$, with $x \in \mathbb{Z}/N\mathbb{Z}$ and $\mathbf{d}$ a prefix. An edge changes either x by one on the cycle or the prefix by one parent-child step. The graph has $N\sum_{r=0}^{h} p^r$ vertices. It includes all tree depths, unlike a cylinder register which contains only the residues at its last depth.

Let $\vec{E}$ denote the directed arcs and let $\tau_2 = 2$, $\tau_p = p - 1$ for odd p. The finite arc Hilbert space is

$$\mathcal{H}_{\text{arc}} = \ell^2(\vec{E}) \otimes \mathbb{C}^{\tau_p}.$$

For each vertex of degree d_v, define the normalized outgoing vector $s_v = d_v^{-1/2}(1, \dots, 1)$ and the Grover coin $G_v = 2|s_v\rangle\langle s_v| - I$. Its square is I, because $|s_v\rangle\langle s_v|$ is a projection. Take any local unitary R_v on the torsion or sign factor and put

$$C = \bigoplus_v (G_v \otimes R_v).$$

The reversal shift S sends the arc $v \to u$ to $u \to v$, retaining the internal register. It is a permutation with $S^* = S$ and $S^2 = I$. The one-tick walk

$$U_{p,h,N} = SC \tag{134}$$

is unitary because $U_{p,h,N}^* U_{p,h,N} = C^*S^*SC = C^*C = I$. Here locality concerns the graph edges, with no continuum limit assumed.

16.2 Local conservation and the product cone

Write $\psi^n(v,u)$ for the internal amplitude on an arc leaving v before the coin, and $\phi^n = C\psi^n$ for its post-coin value. Define

$$\rho_v^n = \sum_{u \sim v} \|\psi^n(v,u)\|^2, \qquad F_{v\to u}^n = \|\phi^n(v,u)\|^2.$$

The local unitary coin gives $\sum_u F^n_{v\to u} = \rho^n_v$. The reversal shift gives $\rho^{n+1}_v = \sum_u F^n_{u\to v}$. Subtraction proves

$$\rho^{n+1}_v - \rho^n_v = \sum_{u\sim v}(F^n_{u\to v} - F^n_{v\to u}). \tag{135}$$

For every set of vertices A, summing this equation cancels every internal directed flux once with its opposite occurrence. Hence

$$\sum_{v\in A}(\rho^{n+1}_v - \rho^n_v) = \sum_{\substack{u\notin A,\ v\in A\\ u\sim v}} F^n_{u\to v} - \sum_{\substack{v\in A,\ u\notin A\\ u\sim v}} F^n_{v\to u}. \tag{136}$$

Taking A to be all vertices gives conservation of total probability. The coarse identity uses a subset of vertices and does not presume amplitude-level Haar refinement.

Proposition 16.1 (Exact finite propagation)**.** *If the initial outgoing amplitudes are supported at v_0, then*

$$d_{C_N}(x, x_0) + d_T(\mathbf{d}, \mathbf{d}_0) > n \quad\Longrightarrow\quad \rho^n_{(x,\mathbf{d})} = 0. \tag{137}$$

For two prefixes, writing lcp *for the length of their longest common prefix,*

$$d_T(\mathbf{d}, \mathbf{d}_0) = |\mathbf{d}| + |\mathbf{d}_0| - 2\,\mathrm{lcp}(\mathbf{d}, \mathbf{d}_0).$$

In particular, each tick allows at most one physical edge or one tree edge, and $\big||\mathbf{d}| - |\mathbf{d}_0|\big| \le n$.

Proof. Every path in a Cartesian product graph uses a number of physical edges at least the physical distance and a number of tree edges at least the tree distance. Their sum is a lower bound on path length. Following a shortest physical path and then a shortest tree path attains that sum, so it is the product graph distance. In a rooted tree the shortest path rises to the longest common prefix and then descends, giving the displayed formula.

The claim at $n = 0$ is the support hypothesis. A coin mixes outgoing arcs only at their common source, so it does not enlarge the vertex support. The shift moves the source across one edge. A state supported in the radius-n ball is therefore supported in the radius-$(n+1)$ ball after a tick. Induction proves the result. Since each graph edge changes only one factor, the individual distance and depth bounds follow. □

The depth difference alone is weaker than the full tree distance. Different prefixes at the same depth can have positive distance, so a plot aggregated over depth does not show the complete graph geometry. The cone is an exact support property of this finite coined model. A Lorentzian continuum dispersion law or a relativistic limiting field equation requires a separate scaling analysis.

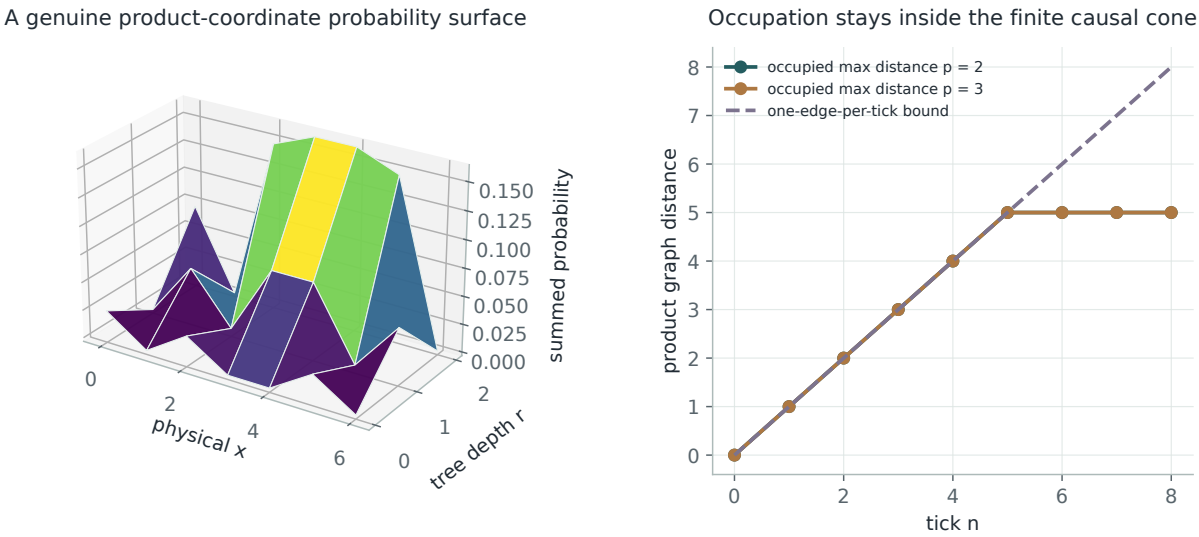


Figure 7. A finite cycle–tree walk, aggregated over vertices at the same tree depth. The graph-distance support check uses the full product graph, including prefix labels. Probability plotted against depth alone displays only an aggregation of that geometry.

16.3 Why growing this tree is not automatically Haar refinement

Example 16.2 (A boundary coin breaks the canonical arc embedding)**.** Take $h \geq 1$ and $N \geq 3$, and embed the old arcs of $C_N \square T_{p,h}$ in $C_N \square T_{p,h+1}$ by identical labels, with zero amplitude on new arcs. Denote this isometry by J. At an old leaf, the old degree is $d = 3$: two cycle neighbors and one parent. The larger graph adds p children, changing the degree to $d + p$. For an old outgoing arc basis vector e_j, the new Grover coin gives amplitude $2/(d + p)$ on each of the p new outgoing child arcs. An internal unitary preserves its internal norm. The subsequent reversal shift sends these amplitudes to new arcs at the children, outside the old arc range. Therefore, for a unit internal state χ,

$$\left\|(I - JJ^*)U_{p,h+1,N}J(e_j \otimes \chi)\right\| = \frac{2\sqrt{p}}{3+p} > 0. \tag{138}$$

The state $JU_{p,h,N}(e_j \otimes \chi)$ has zero amplitude there, so $U_{p,h+1,N}J \neq JU_{p,h,N}$.

Each walk remains unitary, locally conservative, and confined to its product graph cone. A different coin family, embedding, or infinite-tree construction could be investigated. The canonical old-arc inclusion fails the compatibility hypothesis of Theorem 14.1. The cylinder limit and the tree walk are exact models with distinct state spaces.

17 The boundary of the infinite-depth result

The exact result comprises complex L^2 cylinder reconstruction and dynamics for declared real level energies or self-adjoint packet blocks. State convergence uses an L^2 tail; graph norm convergence uses an operator-domain tail. Compact resolvents require energies escaping every bounded set, and finite thermal traces require multiplicity summability. Each claim has a distinct hypothesis.

The source programs check finite restrictions of these formulas. Code 45 uses three towers,

$$(p, M) = (2, 7),\ (3, 5),\ (5, 3),$$

and a fixed random seed. Their terminal zero tails do not certify an infinite domain or compactness; the proofs supply those conclusions. Code 40 tests specified dyadic/ternary interactions and weighted off-diagonal couplings. Code 42 tests finite product walks, whose unitarity does not imply compatibility with the Haar tower.

C15 records finite refinement identities and C16 the conditional infinite-depth construction. Their archived chapter numbers predate the current arrangement. Exact members are chapter 31 for multi-prime dynamics, chapter 33 for product walks, and chapter 36 for the infinite Haar limit. Prime-polyphase semigroups appear in the spectral part. No fresh Lean compilation, unrestricted quantum-field completion, infinite-prime adelic construction, or experimental identification is claimed here.

18 What refinement establishes, and what remains

The finite construction establishes a concrete chain of statements. Unit coordinates factor a finite arithmetic register. Normalized Haar averages identify a coarse Hilbert space inside a fine one. A compatible Hermitian Hamiltonian then carries the same evolution on that range. This last step uses compatibility, rather than normalization alone. When a new coupling breaks the reducing-range property, the hidden sector produces leakage and a nontrivial self-energy.

The spectral construction makes that distinction calculable. Polyphase coordinates display the fine operator's channels without removing any of them. Radial or residue elimination introduces an energy-dependent effective operator on a regular resolvent chart. The memory formulation retains the hidden initial state and the causal history. Chebyshev decimation describes an exact spectral relation where its Dirichlet inverse exists. Hypergeometric companion functions organize entire power-series identities in finite or bounded settings, with additional domains when unbounded operators enter.

The infinite construction is a theorem for a declared hierarchical model. Probability Haar measure and normalized cylinders give a common Hilbert space. Orthogonal depth blocks and self-adjoint weights specify its maximal Hamiltonian domain. Finite compatible blocks then approximate the state evolution strongly. Graph-norm convergence concerns vectors in that domain. Operator-norm convergence and compact resolvent require stronger spectral conditions. A finite walk whose old vertices change degree under extension is outside this compatibility theorem until a suitable comparison map is supplied.

> **The remaining research interfaces**
>
> The analytic adic hierarchy still needs an exact formal translation if it is to carry a Lean certificate. General depth-dependent walks need their own limit argument. A physical fiber theory needs a covariant coupling, a state-selection or production rule, and consistent energy accounting. Each interface is an additional claim with its own hypotheses.

The visual atlas and verification companion make the scope reviewable. Their finite examples test identities, illuminate domain distinctions, and retain negative controls. They do not extrapolate a bounded calculation into an infinite proof. Related archived formal results establish normalized cells, exact tower properties, and dyadic reconstruction statements. They are useful mathematical interfaces, but their theorem names and hypotheses do not certify every result in this volume.

A candidate fiber interpretation must specify how refinement degrees of freedom couple to spacetime and how they are populated. The covariant fiber chapter addresses that question before the cosmological calculations are used. Its energy-budget obstruction remains controlling: downstream transport from a prescribed distribution is meaningful, but it is not evidence that the proposed production history generated that distribution.

Sources and primary context

Project sources

[S1] Komorebi P. *Projector-Refinement Physics, Volume I*. The arithmetic source studies cover unit registers, Haar refinement, finite-adelic dynamics, product walks, radial self-energy, polyphase structure, infinite Haar limits, spectral renormalization, and hypergeometric companion functions.

[S2] Komorebi P. *Volume I Technical Supplement and Independent Review Dossier*. The arithmetic claim dossiers and formula registries specify the scope of the corresponding derivations and computational checks.

[S3] Komorebi P. *Volume I Second Implementation Reproduction*. Finite Haar tests provide internal reproduction evidence for their stated inputs, not for the entire infinite theory.

[S4] Komorebi P. *Komorebi Physics formal source*. The verification appendix identifies related tower and concrete reconstruction results. Shared interfaces are not a certificate for every arithmetic theorem.

Primary context

[C1] NIST Digital Library of Mathematical Functions, §16.2 and §18.5. Generalized hypergeometric series and Chebyshev conventions. Relevant definitions and convergence hypotheses inspected.

[C2] G. Teschl, *Jacobi Operators and Completely Integrable Nonlinear Lattices*, Mathematical Surveys and Monographs 72, American Mathematical Society, 2000. Author's online edition, §§1.1 and 2.1. Jacobi continued-fraction and Weyl-function context. The sign of a Weyl function changes when $(J - z)^{-1}$ is replaced by $(z - J)^{-1}$.

[C3] S. V. Kozyrev, "Wavelet theory as p-adic spectral analysis," *Izvestiya: Mathematics* 66(2), 367–376 (2002). doi:10.1070/IM2002v066n02ABEH000381. Published English Theorems 2 and 7 inspected. The $\mathbb{Q}_p$ Vladimirov setting supplies wavelet context; the present $\mathbb{Z}_p$ level-diagonal construction is proved separately.

[C4] D. Aharonov, A. Ambainis, J. Kempe and U. Vazirani, "Quantum Walks on Graphs," *Proceedings of STOC 2001*, 50–59. doi:10.1145/380752.380758. arXiv:quant-ph/0012090v2, 25 May 2002, §3.1 inspected. Standard coined-walk context; the product-graph arc construction and its continuity identities are proved here.

[C5] P. P. Vaidyanathan, "Multirate digital filters, filter banks, polyphase networks, and applications: a tutorial," *Proceedings of the IEEE* 78(1), 56–93 (1990). doi:10.1109/5.52200. Caltech author-repository record. Metadata and abstract inspected; full PDF retrieval failed. Polyphase history and terminology only.

Scope of the literature check

The hypergeometric bottom parameters used here are positive, so the nonpositive-integer exclusions in §16.2 do not arise. The Chebyshev formulas use the standard first- and second-kind definitions. These scalar conventions do not supply an unbounded-operator domain.

The cited Teschl, Kozyrev, and Aharonov–Ambainis–Kempe–Vazirani passages are used only for the contextual comparisons stated above. Their settings differ from the finite tree, compact adic hierarchy, and irregular arc graph of this volume. The comparison is limited to the explicit uses recorded for each reference. The Vaidyanathan PDF could not be retrieved; its metadata and author abstract establish only the stated historical context. All exact model propositions in this volume carry internal derivations. No exhaustive novelty search or literature-priority claim is made.

PART IV

Matter and cosmological transport

A mathematical refinement fiber becomes a physical proposal only after a coupling and production history are specified. The energy-budget obstruction is examined before the distribution is propagated through cosmological observables.

IN THIS PART

Covariant Refinement-Fiber Dynamics 213

Cosmology, Dark Sectors, and Structure Formation 245

Covariant Refinement-Fiber Dynamics

Symmetry, constraint closure, energy availability, and radiative bounds

Arithmetic refinement supplies a Hilbert-space structure and compatible operators. A physical fiber model additionally needs a covariant action, a state-selection rule, and a consistent energy budget. This account examines those requirements before a refinement distribution is propagated through cosmology.

At finite depth, explicit symmetry assumptions constrain the spectral channel without fixing every coupling. Covariant stress exchange and canonical closure then test compatibility with spacetime dynamics. The production calculation poses a different question: can the proposed crossing actually generate the distribution used by the observable model?

The controlling physical boundary

For the specified benchmark, the unprojected production spectrum exceeds the available crossing energy by $2.6272875374 \times 10^{20}$. A common normalization factor preserves a distribution's shape but does not solve the coupled background and mode equations. The cosmological calculations that follow are conditional on the prescribed normalized distribution, not evidence of a controlled microscopic production history.

The mathematical identities remain meaningful after this physical benchmark fails its energy test. Written theorems, bounded computations, and model assumptions retain separate evidential roles.

CHAPTER 1

Mathematical Structure and Physical Assumptions

1.1 From refinement to a physical fiber

The foundational treatment develops the projector and corner calculus, exact transfer identities, relativistic walks, gauge systems, Lorentzian BV–BFV structures, symmetry-reduced strong-field benchmarks, and adic quantum refinement. The cosmological model takes one proposed scale-covariant matter action into cosmology, kinetic transport, primary CMB templates, nonlinear structure, baryonic response, and linked falsification tests.

These constructions leave an important physical interface. The cosmological action is covariant as written, but the coefficient pattern is postulated, its scalar-fiber contribution was not carried through the full coupled canonical closure, and the production spectrum was normalized after a prescribed-background calculation. The downstream calculations can be internally consistent while the production map that feeds them is not dynamically controlled. A serious assessment must therefore move upstream.

This bridge addresses four questions.

1. Which part of the scale-covariant interaction follows from explicit symmetry assumptions?
2. Does the scalar-fiber sector possess a conserved covariant stress tensor and a closed canonical constraint algebra?
3. Which production identities remain exact for arbitrary refinement order and for a fixed prime tower?
4. Does the specified benchmark lie inside the controlled regime of its rapid-crossing approximation?

1.2 What is already exact

The following ingredients are retained without change.

Object	Status	Boundary
Orthogonal average and detail projectors	theorem	finite or infinite Hilbert-fiber decomposition
Haar detail multiplicity $d_h = (p-1)p^{h-1}$	theorem	chosen prime branching structure
Finite pointwise derivation no-go theorem	theorem	ordinary commutative pointwise product
Matrix-corner derivations $\delta_K X = \mathrm{i}[K, X]$	theorem	noncommutative corner algebra

Alias-free nonlinear gravitational closure	conditional theorem	declared Fourier support headroom
Real Lorentzian BV parent and boundary charge	conditional theorem	regular classical sector and finite resolved domain
Charged spherical mass identity	theorem	spherical canonical model
Infinite-depth Haar martingale and resolvent limits	conditional theorem	stated regularity and operator hypotheses

These results are mathematical statements about their stated models. None of them alone identifies the refinement fiber with a physical hidden sector.

1.3 Assumptions of the matter model

At finite depth H, the cosmological matter model has the schematic form

$$S_H = \int \sqrt{-g}\, \mathrm{d}^4x \left[-\frac{1}{2}(\nabla\phi)^2 - U(\phi) - \frac{1}{2}\langle \nabla\chi, \nabla\chi\rangle_H - \frac{1}{2}\left\langle \chi, \left(m_b^2 + F(\phi)L_H\right)\chi\right\rangle_H \right], \tag{1.1}$$

where

$$L_H = \Lambda^2 \sum_{h=1}^{H} p^{2\sigma h} D_h, \qquad F(\phi) = 1 + \gamma\frac{\phi^2}{v^2}. \tag{1.2}$$

The distinction between structural and fitted input is important.

Quantity	Role in the model	Comment
D_h and d_h	exact	projector and Haar geometry
p, σ	structural input	select branching and spectral order
Λ, m_b, λ, v	microscopic input	not derived from gravity
$F(\phi) = 1 + \gamma\phi^2/v^2$	model proposition	one coefficient across all levels
$\gamma = 16$	benchmark choice	not obtained by matching
$a_\star = 10^{-6}$	benchmark choice	not an invariant physical clock by itself
$\Lambda/A_\star = 10^{-4}$	benchmark choice	fixes the sampling phase of the tower
rapid-crossing Gaussian	conditional calculation	prescribed linear crossing and negligible backreaction
common factor η	energy projection	shape-preserving rescaling, not a solved Bogoliubov backreaction
late transfer and structure	controlled computation	conditional on the normalized distribution

1.4 The decisive benchmark issue

The specified benchmark records

$$\boxed{\eta = 3.8062069178 \times 10^{-21}.} \tag{1.3}$$

Since the projection is defined by $\eta = K_\star/\widetilde{\rho}_\star$, the prescribed-background tower carries

$$\boxed{\frac{\widetilde{\rho}_\star}{K_\star} = \frac{1}{\eta} = 2.6272875374 \times 10^{20}.} \tag{1.4}$$

Interpretation

This is not a small correction to a nearly self-consistent production event. It says that the background used to calculate the spectrum cannot remain approximately unchanged while the spectrum is created. A uniform rescaling preserves convenient ratios, but it is not the solution of the coupled field and mode equations. The production shape must be treated as an ansatz until a backreacting calculation reproduces it.

The late-time Boltzmann and structure calculations are not thereby algebraically invalid. Their interpretation becomes conditional: they show what follows from the prescribed normalized distribution, not what a completed microscopic crossing necessarily produces.

1.5 Mathematical and physical status

The following table distinguishes what is proved, what is computationally tested, and what remains a physical hypothesis.

Statement	Status	Reason
Scale-covariant operator pattern	spectral-channel theorem	after explicit branch assumptions, coefficient still free and an order-zero coupling is allowed
Covariant average-detail exchange	theorem	exact finite-depth Noether identity
Coupled scalar-fiber canonical closure	conditional theorem	total continuum algebra exact, resolved matter H–H bracket exact under headroom
$\sigma = 1$ continuum constants	exact interpolation	not a fixed-prime pointwise limit
Arbitrary-σ constants	theorem	continuous density-of-states model
Specified rapid-crossing spectrum	uncontrolled benchmark ansatz	energy overdraw by $2.6272875374 \times 10^{20}$
Late-time transfer from that spectrum	controlled conditional computation	input distribution remains prescribed
Dark matter and dark energy identification	model proposition	no microscopic production or data likelihood yet

The distinction is decisive for the following cosmological analysis. A conditional distribution can be transported consistently without having been produced by a controlled microscopic solution.

CHAPTER 2

Covariant Finite-Depth Refinement-Fiber Dynamics

2.1 Finite-depth field content

Let $\mathcal{K}_H$ be the detail fiber through depth H,

$$\mathcal{K}_H = \bigoplus_{h=1}^{H} \operatorname{Ran} D_h, \qquad D_h D_j = \delta_{hj} D_h. \tag{2.1}$$

The multiplicity of level h is

$$d_h = (p-1)p^{h-1}. \tag{2.2}$$

The positive refinement operator is

$$L_H = \Lambda^2 \sum_{h=1}^{H} p^{2\sigma h} D_h. \tag{2.3}$$

The average field ϕ is a real spacetime scalar. The detail field χ is a real spacetime scalar valued in $\mathcal{K}_H$. The finite-depth regulator is physical bookkeeping here. It is not sent to infinity until ultraviolet control is supplied.

2.2 Classification of the level operator

Let

$$G_H = \prod_{h=1}^{H} O(d_h) \tag{2.4}$$

act independently within every level. This is the largest real orthogonal symmetry that preserves the decomposition while making no preferred choice inside a degenerate level.

Theorem 2.1 (Level-orthogonal commutant)**.** *If a self-adjoint operator B on $\mathcal{K}_H$ commutes with every element of G_H, then*

$$B = \sum_{h=1}^{H} b_h D_h \tag{2.5}$$

for real coefficients b_h.

Proof. Each $D_h\mathcal{K}_H$ is an irreducible real module for its independent full orthogonal factor. Commutation with the central sign flip on one level removes every off-diagonal map between distinct levels. On one level, commutation with all rotations forces the restriction of B to be a scalar multiple of the identity. Summing the restrictions gives (2.5). □

Refinement is branching rather than an invertible similarity between levels of unequal rank. Choose orthogonal child isometries

$$V_{h,r} : \operatorname{Ran} D_h \longrightarrow \operatorname{Ran} D_{h+1}, \qquad r = 1, \dots, p, \tag{2.6}$$

such that

$$V_{h,r}^* V_{h,r} = D_h, \qquad V_{h,r} V_{h,r}^* \perp V_{h,s} V_{h,s}^*, \qquad \sum_{r=1}^{p} V_{h,r} V_{h,r}^* = D_{h+1}. \tag{2.7}$$

The dimensions agree because $d_{h+1} = p d_h$.

Theorem 2.2 (Branch-homogeneous quadratic operator)**.** *Let $B = B^*$ commute with G_H. If, for every $h < H$ and every child r,*

$$V_{h,r}^* B V_{h,r} = p^{2\sigma} D_h B D_h, \tag{2.8}$$

then

$$B = b_1 \sum_{h=1}^{H} p^{2\sigma(h-1)} D_h. \tag{2.9}$$

If $B \geq 0$, then $b_1 \geq 0$.

Proof. The commutant theorem gives $B = \sum_h b_h D_h$. Pulling the level-$h+1$ restriction back along any child isometry turns (2.8) into $b_{h+1} = p^{2\sigma} b_h$. Hence $b_h = b_1 p^{2\sigma(h-1)}$. A choice $b_1 = \Lambda^2 p^{2\sigma}$ reproduces (2.3). □

The numerical commutant audit used dimensions $1, 2, 4$. The commutant of the independent level algebras had nullity three, one coefficient per level. Adding the two refinement recurrences reduced the nullity to one.

2.3 Leading average-detail interactions

Assume the following EFT conditions.

1. The interaction is local in spacetime.
2. It is quadratic in χ and is separated into fiber orders zero and 2σ.
3. It is invariant under G_H, while the order-2σ part obeys branch homogeneity.
4. The average sector is invariant under $\phi \mapsto -\phi$.
5. The coefficient is analytic near $\phi = 0$.

The general term through fiber order 2σ is

$$-\frac{1}{2} G(\phi) \langle \chi, \chi \rangle_H - \frac{1}{2} F(\phi) \langle \chi, L_H \chi \rangle_H, \qquad \begin{aligned} G(\phi) &= m_b^2 + \kappa \frac{\phi^2}{v^2} + \mathcal{O}(\phi^4/v^4), \\ F(\phi) &= F_0 + \gamma \frac{\phi^2}{v^2} + \mathcal{O}(\phi^4/v^4). \end{aligned} \tag{2.10}$$

After absorbing $F_0 > 0$ into Λ, the leading nontrivial interaction in the order-2σ channel is

$$\boxed{-\frac{\gamma}{2v^2} \phi^2 \langle \chi, L_H \chi \rangle_H .} \tag{2.11}$$

Theorem 2.3 (Spectral-channel classification)**.** *Under conditions 1–5, equation* (2.11) *is the unique nonconstant interaction at quadratic order in both ϕ and χ within the branch-homogeneous order-2σ channel, up to the coefficient γ.*

Proof. The commutant and branch-homogeneity theorems force the order-2σ bilinear to be proportional to L_H. Average-field parity removes the linear term in ϕ. Analyticity leaves a constant followed by ϕ^2/v^2. The constant rescales Λ. The remaining spectral-channel coefficient is γ. □

One-coupling boundary

The symmetry assumptions do not remove the order-zero coupling $\kappa\phi^2 \langle \chi, \chi \rangle_H / (2v^2)$. The cosmological action studied below makes the additional model choice $\kappa = 0$. The theorem therefore establishes one coefficient in the spectral channel, not uniqueness of the complete quadratic EFT. It also does not calculate γ, select p or σ, or derive Λ from gravity.

2.4 Generally covariant action and field equations

Adopting the printed model restriction $\kappa = 0$, the finite-depth covariant system is

$$\boxed{S_H = \int_M \sqrt{-g}\, \mathrm{d}^4x \left[\frac{M_{\mathrm{Pl}}^2}{2}(R - 2\Lambda_{\mathrm{c}}) - \frac{1}{2}\nabla_\mu\phi\nabla^\mu\phi - U(\phi) - \frac{1}{2}\left\langle \nabla_\mu\chi, \nabla^\mu\chi \right\rangle_H - \frac{1}{2}\left\langle \chi, M_H^2(\phi)\chi \right\rangle_H \right] + S_{\mathrm{vis}},} \tag{2.12}$$

where

$$M_H^2(\phi) = m_b^2 + F(\phi)L_H, \qquad F(\phi) = 1 + \gamma\frac{\phi^2}{v^2}. \tag{2.13}$$

The average potential used in the benchmark is

$$U(\phi) = \frac{\lambda}{4}(\phi^2 - v^2)^2, \qquad \mu^2 = \lambda v^2. \tag{2.14}$$

Variation gives

$$\mathcal{E}_\phi := \Box\phi - U'(\phi) - \frac{1}{2}F'(\phi)Y_H = 0, \tag{2.15}$$

$$\mathcal{E}_\chi := \Box\chi - M_H^2(\phi)\chi = 0, \tag{2.16}$$

$$Y_H := \langle \chi, L_H\chi \rangle_H . \tag{2.17}$$

The sign convention uses metric signature $(-,+,+,+)$.

2.5 Stress tensor and exchange identity

The scalar-fiber stress tensor is

$$T^{\mathrm{sf}}_{\mu\nu} = \nabla_\mu\phi\nabla_\nu\phi + \langle \nabla_\mu\chi, \nabla_\nu\chi\rangle_H - g_{\mu\nu}\left[\frac{1}{2}(\nabla\phi)^2 + U(\phi) + \frac{1}{2}\langle\nabla\chi,\nabla\chi\rangle_H + \frac{1}{2}\langle\chi, M_H^2(\phi)\chi\rangle_H\right]. \tag{2.18}$$

Theorem 2.4 (Covariant Ward identity)**.** *For the finite-depth action* (2.12)*,*

$$\boxed{\nabla^\mu T^{\mathrm{sf}}_{\mu\nu} = \mathcal{E}_\phi\nabla_\nu\phi + \langle\mathcal{E}_\chi, \nabla_\nu\chi\rangle_H.} \tag{2.19}$$

The scalar-fiber stress tensor is conserved on shell. A separately diffeomorphism-invariant visible sector adds its own on-shell conserved tensor.

Proof. Differentiate the displayed stress tensor. The kinetic terms reduce to the Euler–Lagrange expressions after commuting covariant derivatives on spacetime scalars. The explicit ϕ derivative of M_H^2 contributes $-\frac{1}{2}F'Y_H\nabla_\nu\phi$. That term combines with the average equation. The remaining detail terms combine with (2.16), which proves (2.19). □

Split $T = T^{(\phi)} + T^{(\chi)}$ by assigning the complete variable mass to the detail sector. Then, on the two field equations,

$$\nabla^\mu T^{(\chi)}_{\mu\nu} = -\frac{1}{2}F'(\phi)Y_H\nabla_\nu\phi, \tag{2.20}$$

$$\nabla^\mu T^{(\phi)}_{\mu\nu} = +\frac{1}{2}F'(\phi)Y_H\nabla_\nu\phi. \tag{2.21}$$

This is an internal transfer. It is not a violation of covariant conservation.

2.6 Exact FLRW exchange

For homogeneous ϕ and detail modes, define

$$\rho_\phi = \frac{1}{2}\dot\phi^2 + U, \qquad p_\phi = \frac{1}{2}\dot\phi^2 - U, \tag{2.22}$$

$$\rho_\chi = \frac{1}{2}\langle\dot\chi,\dot\chi\rangle_H + \frac{1}{2}\langle\chi, M_H^2\chi\rangle_H, \qquad p_\chi = \frac{1}{2}\langle\dot\chi,\dot\chi\rangle_H - \frac{1}{2}\langle\chi, M_H^2\chi\rangle_H. \tag{2.23}$$

The field equations imply

$$\boxed{\dot\rho_\chi + 3H(\rho_\chi + p_\chi) = \frac{1}{2}\dot F\,Y_H,} \tag{2.24}$$

$$\boxed{\dot\rho_\phi + 3H(\rho_\phi + p_\phi) = -\frac{1}{2}\dot F\,Y_H.} \tag{2.25}$$

At the locked origin, $\phi = 0$ and $\dot F = 0$. Separate dust-like detail evolution and constant average potential are then compatible. Once release begins, forcing $\rho_\chi \propto a^{-3}$ while holding the average potential fixed generally violates (2.24)–(2.25).

The random finite-depth Ward check returned a total residual of zero at recorded precision. Omitting the average exchange term left the nonzero residual

$$9.7207340475. \tag{2.26}$$

2.7 Tree-level stability

Theorem 2.5 (Finite-depth stability conditions)**.** *Suppose the kinetic terms in* (2.12) *have their displayed signs. On a field domain $\mathcal{U}$, the scalar-fiber sector has no ghost, gradient, or detail tachyon if*

$$F(\phi) > 0 \quad \textit{and} \quad m_b^2 + F(\phi)\Lambda^2 p^{2\sigma} > 0 \qquad (\phi \in \mathcal{U}). \tag{2.27}$$

At the locked origin, local average stability additionally requires

$$\boxed{M_\phi^2(0) = -\mu^2 + \frac{\gamma}{v^2} Y_H > 0.} \tag{2.28}$$

Proof. The principal symbol is diagonal and canonical, so ghosts and gradient instabilities are absent. The smallest detail eigenvalue occurs at $h = 1$, which gives the second condition in (2.27). Differentiating the potential twice at $\phi = 0$ gives (2.28). □

For the polynomial truncation $F = 1 + \gamma\phi^2/v^2$, global boundedness without higher stabilizing operators requires $\gamma \geq 0$ when L_H has a positive eigenvalue. If $\gamma < 0$, sufficiently large $|\phi|$ makes F negative and the potential is unbounded along a detail eigenvector. This statement concerns the truncated EFT. A positive nonpolynomial completion can change the global conclusion.

2.8 Polynomial Lorentzian first-order parent

The metric action can be embedded into the real polynomial tetrad parent used by the foundational volume. Let e^I be a tetrad, let

$$\Sigma_I = \frac{1}{3!}\epsilon_{IJKL}e^J \wedge e^K \wedge e^L, \qquad \mathrm{vol}_e = \frac{1}{4!}\epsilon_{IJKL}e^I \wedge e^J \wedge e^K \wedge e^L, \tag{2.29}$$

and introduce an auxiliary Lorentz vector u^I and a fiber-valued Lorentz vector w^I. Define

$$\boxed{\begin{aligned} S^{(1)}_{\mathrm{sf}} = \int_M \Bigg[& \frac{1}{2}\left(u_I u^I + \left\langle w_I, w^I \right\rangle_H\right)\mathrm{vol}_e - u^I\, \mathrm{d}\phi \wedge \Sigma_I \\ & - \left\langle w^I,\, \mathrm{d}\chi \right\rangle_H \wedge \Sigma_I - V_H(\phi,\chi)\,\mathrm{vol}_e \Bigg], \end{aligned}} \tag{2.30}$$

where

$$V_H = U(\phi) + \frac{1}{2}\left\langle \chi, M_H^2(\phi)\chi \right\rangle_H . \tag{2.31}$$

Theorem 2.6 (First-order equivalence)**.** *On a nondegenerate tetrad chart, eliminating u^I and w^I from (2.30) gives the scalar-fiber matter part of (2.12).*

Proof. Variation with respect to the auxiliaries gives

$$u_I = e_I{}^\mu \partial_\mu \phi, \qquad w_I = e_I{}^\mu \partial_\mu \chi. \tag{2.32}$$

Substitution completes the square and yields the two canonical kinetic terms with the potential V_H. □

The parent is real and polynomial when F and U are truncated polynomials. The fields ϕ and χ are Lorentz scalars. The auxiliaries are Lorentz vectors. Diffeomorphisms act by Lie derivatives, and the internal level group commutes with spacetime gauge transformations. The off-shell diffeomorphism–Lorentz algebra is therefore the direct scalar-matter extension of the existing parent algebra. Classical BRST nilpotency follows with the ordinary scalar Lie action. A quantum BV measure and anomaly calculation remain open.

CHAPTER 3

Scalar-Fiber Canonical and Resolved Closure

3.1 ADM variables

Write the spacetime metric as

$$\mathrm{d}s^2 = -N^2\,\mathrm{d}t^2 + q_{ij}(\,\mathrm{d}x^i + N^i\,\mathrm{d}t)(\,\mathrm{d}x^j + N^j\,\mathrm{d}t). \tag{3.1}$$

Let π_ϕ and Π be the momenta conjugate to ϕ and χ. The scalar-fiber matter constraints are

$$\boxed{\begin{aligned}\mathcal{H}_\perp^{\mathrm{sf}} = \frac{1}{2\sqrt{q}}\left(\pi_\phi^2 + \langle \Pi, \Pi\rangle_H\right) \\ + \frac{\sqrt{q}}{2}\left[q^{ij}\partial_i\phi\partial_j\phi + q^{ij}\left\langle \partial_i\chi, \partial_j\chi\right\rangle_H + 2V_H(\phi,\chi)\right],\end{aligned}} \tag{3.2}$$

$$\boxed{\mathcal{H}_i^{\mathrm{sf}} = \pi_\phi\partial_i\phi + \langle \Pi, \partial_i\chi\rangle_H\,.} \tag{3.3}$$

The smeared generators are

$$H_{\mathrm{sf}}[N] = \int_\Sigma N\mathcal{H}_\perp^{\mathrm{sf}}\,\mathrm{d}^3x, \qquad D_{\mathrm{sf}}[\boldsymbol{N}] = \int_\Sigma N^i\mathcal{H}_i^{\mathrm{sf}}\,\mathrm{d}^3x. \tag{3.4}$$

3.2 Matter brackets and their metric term

Theorem 3.1 (Matter momentum and Hamiltonian brackets)**.** *On a smooth spatial chart with boundary conditions that remove surface terms, the constraints* (3.2)–(3.3) *satisfy*

$$\{D_{\mathrm{sf}}[\boldsymbol{N}], D_{\mathrm{sf}}[\boldsymbol{M}]\} = D_{\mathrm{sf}}[[\boldsymbol{N}, \boldsymbol{M}]], \tag{3.5}$$

$$\{H_{\mathrm{sf}}[N], H_{\mathrm{sf}}[M]\} = D_{\mathrm{sf}}[\boldsymbol{S}(N, M)], \tag{3.6}$$

where

$$S^i(N, M) = q^{ij}(N\partial_j M - M\partial_j N). \tag{3.7}$$

The result holds for every differentiable scalar potential V_H, including the scale-covariant fiber interaction. The pure-matter mixed bracket is instead

$$\boxed{\begin{aligned}\{D_{\mathrm{sf}}[\boldsymbol{N}], H_{\mathrm{sf}}[M]\} = H_{\mathrm{sf}}[\mathcal{L}_{\boldsymbol{N}} M] \\ + \int_\Sigma \frac{\delta H_{\mathrm{sf}}[M]}{\delta q_{ij}}\mathcal{L}_{\boldsymbol{N}} q_{ij}\,\mathrm{d}^3x.\end{aligned}} \tag{3.8}$$

Proof. The momentum generator acts on ϕ and χ by their scalar Lie derivatives and on the momenta by the corresponding density Lie derivatives. This gives (3.5). It does not act on the metric appearing in H_{sf}. Restoring that missing metric transformation gives the second term in (3.8).

For the Hamiltonian bracket, vary the kinetic and spatial-gradient terms. Terms containing $\partial V_H/\partial\phi$ and the fiber gradient $\partial V_H/\partial\chi$ appear multiplied by NM in both bracket orders and cancel by antisymmetry. The remaining cross terms are

$$\int_\Sigma (N\partial_i M - M\partial_i N)q^{ij}\left(\pi_\phi\partial_j\phi + \langle\Pi, \partial_j\chi\rangle_H\right)\,\mathrm{d}^3x, \tag{3.9}$$

which is (3.6). □

The potential independence of (3.6) is important. Canonical closure tests covariance of the matter action. It does not determine the numerical value of γ or validate the production state.

3.3 Closure after coupling to gravity

Let $H_{\mathrm{G}}, D_{\mathrm{G}}$ denote the gravitational constraints on a regular ADM chart. The total generators are

$$H_{\mathrm{tot}} = H_{\mathrm{G}} + H_{\mathrm{sf}}, \qquad D_{\mathrm{tot}} = D_{\mathrm{G}} + D_{\mathrm{sf}}. \tag{3.10}$$

Theorem 3.2 (Total hypersurface deformation)**.** *On every regular ADM chart where the gravitational bracket is defined, minimal coupling gives*

$$\{D_{\mathrm{tot}}[\boldsymbol{N}], D_{\mathrm{tot}}[\boldsymbol{M}]\} = D_{\mathrm{tot}}[[\boldsymbol{N}, \boldsymbol{M}]], \tag{3.11}$$
$$\{D_{\mathrm{tot}}[\boldsymbol{N}], H_{\mathrm{tot}}[M]\} = H_{\mathrm{tot}}[\mathcal{L}_{\boldsymbol{N}} M], \tag{3.12}$$
$$\{H_{\mathrm{tot}}[N], H_{\mathrm{tot}}[M]\} = D_{\mathrm{tot}}[\boldsymbol{S}(N, M)]. \tag{3.13}$$

Proof. Minimal coupling makes $\mathcal{H}_\perp^{\mathrm{sf}}$ a scalar density of weight one under the total momentum generator. The metric-variation term in (3.8) is supplied by the gravitational action on q_{ij}. The mixed metric variations in the total Hamiltonian bracket combine with the two pure brackets, yielding the displayed algebra. □

> **Closure boundary**
>
> This result is a classical constraint theorem. It is not a quantum anomaly theorem. It also does not turn a nonlocal spectral derivative into a strictly finite-range microscopic lattice action.

3.4 Alias-free resolved realization

On the periodic line, let $\mathcal{T}_K$ denote trigonometric polynomials of band at most K. Sampling on N_x points is injective below Nyquist. If every product generated in one bracket remains below $N_x/2$, the sampled derivative, product, and normalized trace coincide with their continuum values. The proof of the matter H–H identity (3.6) then transfers term by term. The resolved test below does not test the mixed bracket or the dynamical metric sector.

For the one-dimensional scalar-fiber certificate, the Hamiltonian density is

$$\mathcal{H}^{\mathrm{sf}} = \frac{1}{2}\left(\pi_\phi^2 + \langle \Pi, \Pi \rangle_H\right) + \frac{1}{2}\left((\delta\phi)^2 + \langle \delta\chi, \delta\chi \rangle_H\right) + V_H. \tag{3.14}$$

The exact target identity is

$$\{H_N[n], H_N[m]\} = D_N[n\delta m - m\delta n]. \tag{3.15}$$

Proposition 3.3 (Resolved matter Hamiltonian bracket)**.** *If the fields and smearings lie in a band for which every product in the variation and bracket is resolved without aliasing, equation* (3.15) *holds exactly in the sampled algebra.*

Proof. The alias-free product lemma identifies every sampled multiplication and derivative in the bracket with the corresponding trigonometric-polynomial operation. The continuum proof of (3.6) therefore applies without a truncation remainder. □

The deterministic test used $N_x = 128$ and band $K = 3$. Across twenty random trials, the maximum relative residual was

$$\boxed{\epsilon_{\mathrm{resolved}} = 9.907025 \times 10^{-16}.} \tag{3.16}$$

3.5 Negative controls

Three changes break the hypotheses in distinguishable ways.

1. Raising the band to $K = 18$ on $N_x = 64$ aliases products. The residual becomes 1.8994349731.
2. Replacing the spectral derivative by a centered difference destroys the exact product rule. The residual becomes 0.0357793042.
3. Removing the detail spatial-gradient term while retaining its momentum contribution makes the Hamiltonian and momentum generators inconsistent. The residual becomes 1.2401043846.

These tests are stronger than substituting one favorable trajectory. They show that the small residual is tied to the theorem's structural assumptions.

3.6 Relation to the polynomial Lorentzian parent

The first-order matter action (2.30) can be appended to the existing Einstein–Cartan parent before the $3 + 1$ split. The normal tetrad multiplier then receives one additional scalar-fiber constraint, while tangential diffeomorphisms receive (3.3) after regular reduction. Level rotations commute with Lorentz transformations and spacetime diffeomorphisms.

At finite spectral resolution, the complete polynomial BV master identity transfers whenever every monomial generated by the action and antibracket lies below Nyquist. A conservative headroom rule can be obtained by counting the eight-field potential monomial together with antifield variations. This rule is sufficient rather than sharp. Strict finite-range locality, a quantum BV measure, and uniform continuum estimates remain separate tasks.

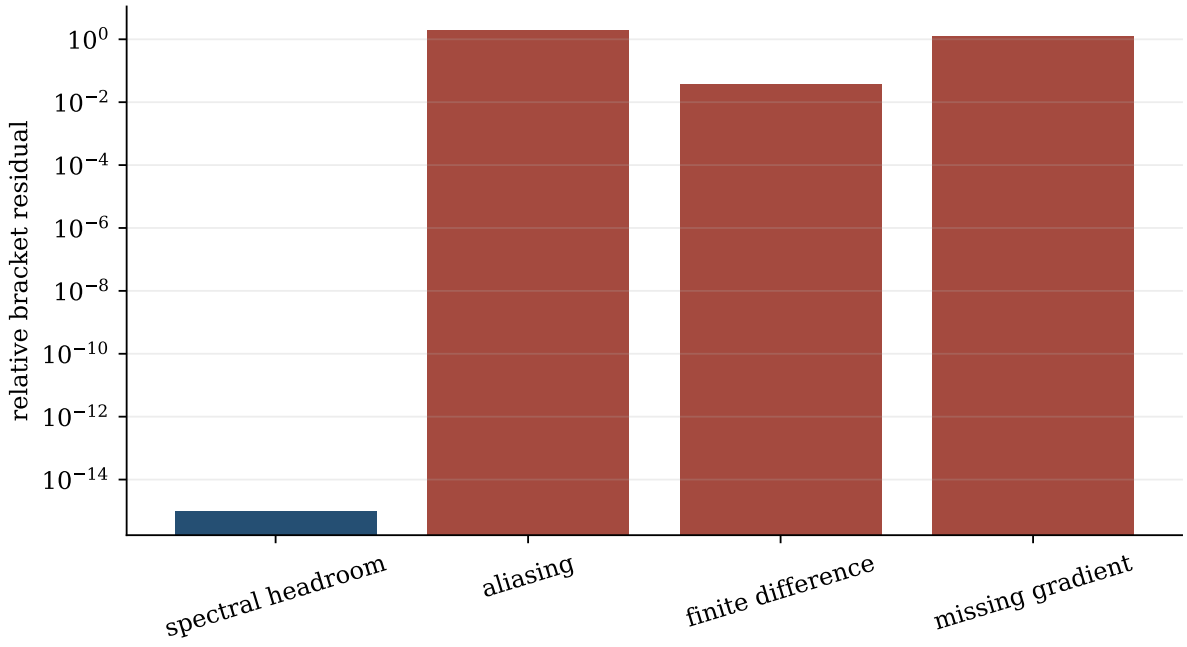


Figure 3.1. The scalar-fiber Hamiltonian bracket closes in the resolved spectral regime. Aliasing, a derivative without an exact product rule, and an incomplete detail constraint all produce large defects.

3.7 What has and has not been closed

Statement	Status	Boundary
finite-depth covariant scalar-fiber action	closed	classical EFT
total stress conservation	closed	on shell
average-detail exchange law	closed	exact at finite depth
matter H–H bracket and total ADM algebra	closed	regular classical chart with gravity included
resolved matter Hamiltonian bracket	closed conditional	alias-free spectral headroom
polynomial first-order parent	closed	nondegenerate tetrad chart
strict finite-range nonlinear realization	open	corner-to-continuum construction required
quantum master equation and anomaly	open	regulator and measure required
self-consistent production state	open	addressed diagnostically in Chapter 5

CHAPTER 4

Arbitrary-Order Production and Fixed-Prime Sampling

4.1 Continuous density-of-states interpolation

Let

$$\alpha := \frac{1}{\sigma}, \qquad m_h \simeq \Lambda p^{\sigma h}, \qquad N_h = p^h - 1. \tag{4.1}$$

Interpolating the endpoint counting function gives

$$N_{\mathrm{c}}(m) = \left(\frac{m}{\Lambda}\right)^{\alpha} - 1, \qquad \frac{\mathrm{d}N_{\mathrm{c}}}{\mathrm{d}m} = \frac{1}{\sigma\Lambda^{\alpha}} m^{\alpha-1}. \tag{4.2}$$

This defines a continuous model. It is exact as an interpolation of the endpoint law. Section 4.4 shows why it is not an ordinary pointwise limit of a fixed-prime tower.

At a prescribed linear crossing, take

$$\nu(m) = Am, \qquad f(k,m) = \eta \exp\left[-\pi\left(\frac{m}{A} + \frac{k^2}{Am}\right)\right]. \tag{4.3}$$

The Gaussian momentum integral gives the number spectrum

$$\boxed{\frac{\mathrm{d}n}{\mathrm{d}m} = \frac{\eta A^{3/2}}{8\pi^3\sigma\Lambda^{\alpha}} m^{\alpha+1/2} \mathrm{e}^{-\pi m/A}.} \tag{4.4}$$

Multiplication by m gives the captured rest-energy spectrum

$$\boxed{\frac{\mathrm{d}\rho_R}{\mathrm{d}m} = \frac{\eta A^{3/2}}{8\pi^3\sigma\Lambda^{\alpha}} m^{\alpha+3/2} \mathrm{e}^{-\pi m/A}.} \tag{4.5}$$

4.2 Exact arbitrary-order constants

Theorem 4.1 (Continuous arbitrary-order production). *For $\sigma > 0$, the continuous model (4.2)–(4.3) has*

$$n_\star = C_N(\sigma)\eta \frac{A^{\alpha+3}}{\Lambda^{\alpha}}, \qquad C_N(\sigma) = \frac{\Gamma(\alpha+\frac{3}{2})}{8\sigma\pi^{\alpha+9/2}}, \tag{4.6}$$

$$\rho_{R,\star} = C_R(\sigma)\eta \frac{A^{\alpha+4}}{\Lambda^{\alpha}}, \qquad C_R(\sigma) = \frac{\Gamma(\alpha+\frac{5}{2})}{8\sigma\pi^{\alpha+11/2}}, \tag{4.7}$$

$$\rho_\star = C_E(\sigma)\eta \frac{A^{\alpha+4}}{\Lambda^{\alpha}}, \qquad C_E(\sigma) = \frac{\Gamma(\alpha+4)B(\frac{3}{2},\alpha+2)}{4\sigma\pi^{\alpha+6}}, \tag{4.8}$$

$$p_\star = C_P(\sigma)\eta \frac{A^{\alpha+4}}{\Lambda^{\alpha}}, \qquad C_P(\sigma) = \frac{C_E(\sigma)}{2\alpha+7}. \tag{4.9}$$

Consequently,

$$\boxed{w_\star(\sigma) = \frac{1}{7 + 2/\sigma},} \tag{4.10}$$

$$\boxed{f_R(\sigma) = \frac{\rho_{R,\star}}{\rho_\star} = \frac{C_R(\sigma)}{C_E(\sigma)},} \tag{4.11}$$

and

$$\boxed{\frac{\rho_{R,\star}}{n_\star} = m^{(R)}_{\text{peak}} = \frac{A}{\pi}\left(\frac{1}{\sigma} + \frac{3}{2}\right).} \tag{4.12}$$

Proof. Equations (4.6) and (4.7) follow by integrating (4.4) and (4.5) with the gamma integral.

For total energy and pressure, set $k = m \sinh\theta$. The production exponent becomes

$$\pi \frac{m}{A} \cosh^2\theta. \tag{4.13}$$

After the mass integral, the energy angular factor is

$$I_E = \int_0^\infty \sinh^2\theta \cosh^{-2\alpha-6}\theta \, \mathrm{d}\theta = \frac{1}{2} B\left(\frac{3}{2}, \alpha + 2\right). \tag{4.14}$$

The pressure angular factor is

$$I_P = \frac{1}{3}\int_0^\infty \sinh^4\theta \cosh^{-2\alpha-8}\theta \, \mathrm{d}\theta = \frac{1}{6} B\left(\frac{5}{2}, \alpha + 2\right). \tag{4.15}$$

Their ratio is

$$\frac{I_P}{I_E} = \frac{1}{3}\frac{B(\frac{5}{2}, \alpha+2)}{B(\frac{3}{2}, \alpha+2)} = \frac{1}{2\alpha + 7}, \tag{4.16}$$

which proves (4.10). The gamma recurrence gives the mean rest mass, and differentiating (4.5) gives the same value for its peak. □

At $\sigma = 1$, these formulas reduce to

$$w_\star = \frac{1}{9}, \qquad f_R = \frac{525\pi}{2048}, \qquad m^{(R)}_{\text{peak}} = \frac{5A}{2\pi}. \tag{4.17}$$

4.3 Prime independence of the interpolation

The prime p cancels from (4.2) because

$$p^h = \left(\frac{m_h}{\Lambda}\right)^{1/\sigma}. \tag{4.18}$$

This gives a continuous universality class labeled by σ. Prime independence at this stage does not mean that the discrete tower forgets p. The spacing in $\log m$ is $\sigma \log p$, which remains finite when Λ/A is reduced.

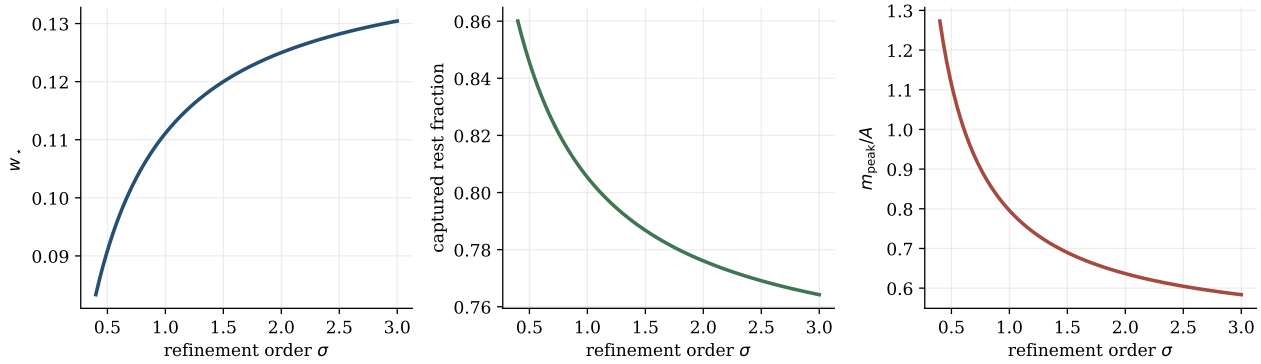


Figure 4.1. The continuous arbitrary-order family. Increasing σ raises the production equation of state while lowering the captured rest fraction and the selected mass in units of A.

4.4 Fixed-prime Mellin correction

Let $F(y)$ be a single-state dimensionless observable with $y = m/A$, and define

$$x = \frac{\Lambda}{A}, \qquad c_p = \frac{p-1}{p}, \qquad a = \sigma \log p. \tag{4.19}$$

Ignoring a lower-end correction that vanishes as $x \to 0$ for the production kernels, extend the level sum bilaterally:

$$S_F(x) = c_p \sum_{h \in \mathbb{Z}} p^h F(x p^{\sigma h}). \tag{4.20}$$

Let

$$\mathcal{M}F(s) = \int_0^\infty y^{s-1} F(y)\, \mathrm{d}y. \tag{4.21}$$

Theorem 4.2 (Log-periodic fixed-prime decomposition)**.** *Suppose $F(y) = \mathcal{O}(y^\beta)$ as $y \to 0$ for some $\beta > 0$, decays sufficiently rapidly at infinity, and*

$$g_x(t) = c_p \mathrm{e}^{t \log p} F(x \mathrm{e}^{at}) \tag{4.22}$$

together with its Fourier transform satisfies the integrability and summability hypotheses of Poisson summation. A sufficient stronger condition is that g_x is Schwartz. Then

$$\boxed{S_F(x) = \frac{c_p}{a} x^{-\alpha} \sum_{k \in \mathbb{Z}} x^{\mathrm{i}\omega_k} \mathcal{M}F(\alpha - \mathrm{i}\omega_k), \qquad \omega_k = \frac{2\pi k}{a}.} \tag{4.23}$$

Proof. Apply Poisson summation to

$$g(t) = c_p \mathrm{e}^{t \log p} F(x \mathrm{e}^{at}). \tag{4.24}$$

In the Fourier transform, substitute $y = x\mathrm{e}^{at}$. The factor $\mathrm{e}^{-2\pi \mathrm{i} k t}$ becomes $(y/x)^{-\mathrm{i}\omega_k}$, and the remaining integral is the Mellin transform in (4.23). The small-y hypothesis also bounds the omitted physical levels $h \le 0$. Their absolute contribution is $\mathcal{O}(x^\beta)$ and is negligible relative to the $x^{-\alpha}$ leading abundance. □

The $k = 0$ term has coefficient

$$\frac{c_p}{\sigma \log p} x^{-\alpha} \mathcal{M}F(\alpha). \tag{4.25}$$

The derivative interpolation (4.2) instead has coefficient

$$\frac{1}{\sigma} x^{-\alpha} \mathcal{M}F(\alpha). \tag{4.26}$$

Their mean amplitude ratio is

$$\boxed{\frac{c_p}{\log p} = \frac{p-1}{p \log p}.} \tag{4.27}$$

It approaches one only in a formal dense-branching limit $p \to 1$, not along the physical prime sequence. The $k \neq 0$ terms add a periodic dependence on $\log x$.

Corollary 4.3 (No fixed-prime convergence to the derivative interpolation)**.** *After the natural x^{α} rescaling, reducing Λ/A at fixed p translates the sampling grid in $\log m$ but does not shrink its spacing. The rescaled discrete normalization therefore retains a prime-dependent mean offset and log-periodic corrections rather than converging to the derivative interpolation. Ratios of observables can nevertheless be much more accurate because common amplitude factors cancel.*

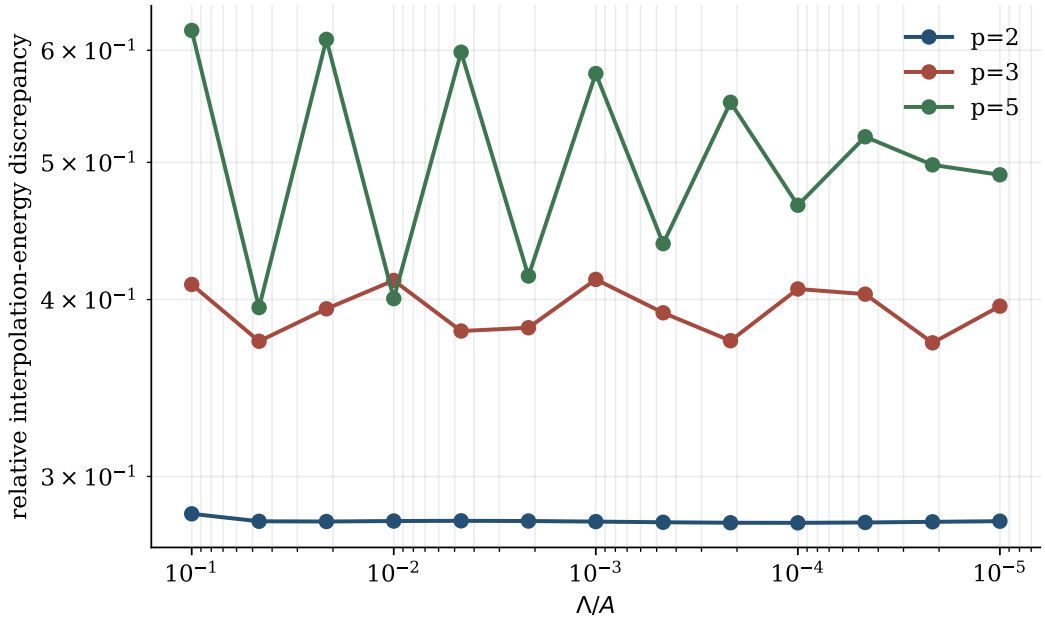


Figure 4.2. Relative discrepancy between the discrete total production energy and the derivative interpolation at $\sigma = 1$. Decreasing Λ/A does not remove the fixed logarithmic spacing. The nearly constant dyadic offset and the larger prime-dependent oscillations agree with the Mellin decomposition.

4.5 Correction to the interpretation of the continuum formulas

The continuous constants remain exact for the density-of-states model defined by (4.2). They also organize robust ratios in the dense mass band. They should not be described as a convergent fixed-prime limit without an additional averaging operation, a refinement in logarithmic spacing, or a theorem that suppresses the nonzero Mellin harmonics.

This distinction matters most for absolute abundance. The cosmological benchmark uses the direct discrete tower calculation when imposing its energy projection, so its discrete normalization is not replaced by the continuous coefficient. The estimate explains why the reported dyadic ratios can be close to the continuous values even though the absolute energy coefficient differs substantially.

CHAPTER 5

Production Backreaction and the Frozen Benchmark

5.1 Control conditions for the rapid-crossing spectrum

The Gaussian occupation (4.3) follows from a prescribed approximately linear crossing. Three independent small parameters are required near the populated band:

$$\epsilon_H := \frac{H_\star^2}{A_\star m_{\rm peak}} \ll 1, \tag{5.1}$$

$$\epsilon_{\rm lin} := \frac{|\ddot{\phi}_\star|}{|\dot{\phi}_\star|\sqrt{A_\star m_{\rm peak}}} \ll 1, \tag{5.2}$$

$$\epsilon_{\rm br} := \frac{\widetilde{\rho}_\star}{K_\star} \ll 1. \tag{5.3}$$

The first two control expansion and curvature across the event. The third controls depletion by the produced modes. A posteriori normalization cannot replace (5.3) because depletion changes the background that generated the Bogoliubov coefficients.

5.2 Benchmark energy-overdraw theorem

The benchmark with $\lambda = 1$ defines a shape-preserving projection

$$f_h(k) \longmapsto \eta f_h(k), \qquad \eta = \min\left(1, \frac{K_\star}{\widetilde{\rho}_\star}\right). \tag{5.4}$$

Theorem 5.1 (Failure of negligible backreaction)**.** *For the specified benchmark record,*

$$\eta = 3.8062069178 \times 10^{-21} \tag{5.5}$$

and therefore

$$\boxed{\epsilon_{\rm br} = \frac{1}{\eta} = 2.6272875374 \times 10^{20}.} \tag{5.6}$$

The benchmark is outside the negligible-backreaction domain by more than twenty orders of magnitude.

Proof. The value of η is read from the frozen discrete-tower calculation. On the active branch of (5.4), $\eta = K_\star/\widetilde{\rho}_\star$. Inversion gives (5.6). □

Physical consequence

The exact ratios $w_\star = 1/9$ and $f_R = 525\pi/2048$ are properties of the projected continuous ansatz at $\sigma = 1$. They are not predictions of the coupled crossing until a backreacting solution reproduces the same spectral shape. With $\epsilon_{\rm br} \gg 1$, strong depletion is mandatory.

5.3 Hubble-loss correction before the crossing

If the average field begins at rest at ϕ_i and reaches $\phi = 0$, the mechanical crossing energy in FLRW obeys

$$\boxed{K_{\star^-} = U(\phi_i) - U(0) - 3\int_{t_i}^{t_\star} H\dot{\phi}^2 \, \mathrm{d}t - E_{\rm other},} \tag{5.7}$$

where $E_{\rm other}$ denotes the integrated work removed from the average field before the crossing. It need not equal the final redshifted energy density of the receiving sector. The formula without the integral is exact only in a negligible-expansion and negligible-transfer limit. The benchmark does not fix the absolute ratio $\mu/H_\star$, so that approximation cannot be checked from the dimensionless record alone.

5.4 Continuous order-one backreaction estimate

In the arbitrary-order continuous model,

$$\widetilde{\rho}_\star = C_E(\sigma)\frac{A^{\alpha+4}}{\Lambda^\alpha}, \qquad K_\star = \frac{v^2A^2}{2\gamma}, \qquad \gamma > 0. \tag{5.8}$$

Introduce

$$a = \frac{A}{\mu}, \qquad \ell = \frac{\Lambda}{\mu}, \qquad \frac{v^2}{\mu^2} = \frac{1}{\lambda}. \tag{5.9}$$

The raw continuous overdraw is

$$R_{\rm c}(a) = 2\gamma\lambda C_E(\sigma)\frac{a^{\alpha+2}}{\ell^\alpha}. \tag{5.10}$$

For the frozen dimensionless input, $R_{{\rm c},in} = 3.6399161857 \times 10^{20}$. If changing A were the only response, the order-one boundary $R_{\rm c} = 1$ would satisfy

$$\boxed{\frac{A_{\rm br,c}}{A_{\rm in}} = R_{{\rm c},in}^{-1/(\alpha+2)}.} \tag{5.11}$$

For the benchmark, $\alpha = 1$ and

$$\frac{A_{\rm br,c}}{A_{\rm in}} = 1.40056372 \times 10^{-7}, \tag{5.12}$$

$$\frac{A_{\rm br,c}}{\mu} = 271.492078, \tag{5.13}$$

$$\frac{\Lambda}{A_{\rm br,c}} = 713.998218. \tag{5.14}$$

The lightest level would then have suppression exponent

$$\pi \frac{m_1}{A_{\mathrm{br,c}}} \simeq 4486.183112. \tag{5.15}$$

The continuous populated-band approximation is self-inconsistent at that point. This is not a paradox. It says that slowing the same continuous spectrum cannot reach its order-one boundary while holding the microscopic tower fixed. The raw continuous coefficient is used here rather than the fixed-prime projection factor η.

5.5 Direct discrete order-one threshold

The next calculation keeps the benchmark's fixed Λ, p, σ, γ, v, and Gaussian single-crossing formula. It changes only A and evaluates the cutoff-controlled discrete momentum moments. Define

$$R_{\mathrm{disc}}(A) := \frac{\widetilde{\rho}_{\mathrm{disc}}(A)}{v^2 A^2/(2\gamma)}. \tag{5.16}$$

Controlled diagnostic 5.2 (Discrete backreaction threshold)**.** The equation $R_{\mathrm{disc}}(A) = 1$ has the numerical solution

$$\boxed{\frac{A_{\mathrm{br,d}}}{\mu} = 49212.192097, \qquad \frac{\Lambda}{A_{\mathrm{br,d}}} = 3.938960080.} \tag{5.17}$$

At this point,

$$\boxed{w_{\mathrm{br,d}} = 0.0187227533, \qquad f_{R,\mathrm{br,d}} = 0.9712627819, \qquad \frac{\langle m \rangle_R}{A} = 7.8779201618.} \tag{5.18}$$

The calculation uses 128-node generalized Gauss–Laguerre momentum quadrature and the exact discrete degeneracies. Three levels fall below the declared suppression cutoff. The first carries essentially all of the production weight.

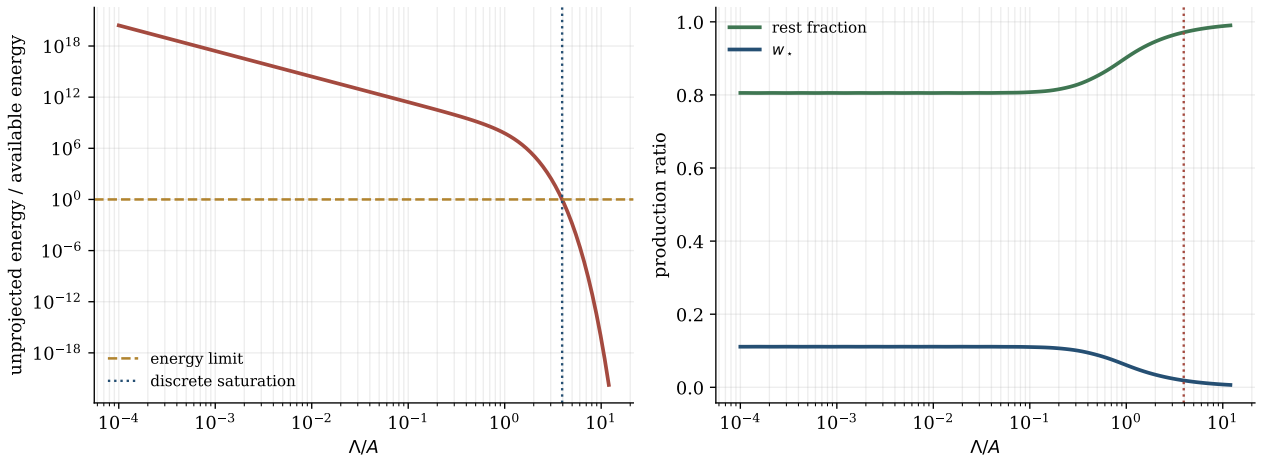


Figure 5.1. Left: the unprojected energy relative to the crossing reservoir as Λ/A is increased at fixed microscopic Λ. The order-one backreaction threshold occurs after the continuous populated-band regime has been left. Right: the production state becomes more rest dominated and its pressure falls as the tower sparsifies.

Diagnostic boundary

Equation (5.17) is not a time-dependent backreaction solution and is not an energy-partition equation for the original crossing. It marks where the prescribed spectrum and kinetic reservoir become comparable inside a one-parameter Gaussian family. Its value is useful because the family has already become sparse and cold at that boundary. A dynamical solution can depart further.

5.6 Interaction-strength obstruction

For a detail eigenvalue $\lambda_h = \Lambda^2 p^{2\sigma h}$, the canonical four-dimensional action contains the quartic interaction

$$-\frac{1}{2} g_h \phi^2 \chi_h^2, \qquad g_h = \gamma \frac{\lambda_h}{v^2}. \tag{5.19}$$

Using $v^2 = \mu^2/\lambda$ and $m_h^2 \simeq \lambda_h$ gives

$$g_h \simeq \gamma \lambda \left(\frac{m_h}{\mu} \right)^2. \tag{5.20}$$

At the frozen discrete peak with $\gamma = 16$ and $\lambda = 1$,

$$\boxed{g_{\text{peak}} \simeq 4.03467400 \times 10^{19}.} \tag{5.21}$$

A nominal loop-control condition $g_{\text{peak}} < 1$ would instead require

$$\lambda < 2.47851499 \times 10^{-20} \tag{5.22}$$

with the other benchmark ratios held fixed. At $\lambda = 1$, even the lightest level at the discrete threshold has

$$g_1 \simeq 2.40485311 \times 10^{12}. \tag{5.23}$$

At fixed calibrated ratios A/μ and Λ/μ, the two diagnostics are parametrically linked:

$$R_{\text{br}}(\lambda) = \frac{\lambda}{\lambda_{\text{frozen}}} R_{\text{br}}(\lambda_{\text{frozen}}). \tag{5.24}$$

For $\lambda_{\text{frozen}} = 1$, avoiding energy overdraw requires

$$\boxed{\lambda \leq 3.8062069178 \times 10^{-21}.} \tag{5.25}$$

At that boundary, the peak and lightest-level couplings would be

$$g_{\text{peak}} = 0.1535680410, \qquad g_1 = 9.15336853 \times 10^{-9}. \tag{5.26}$$

Perturbative boundary

The printed $\lambda = 1$ benchmark is neither weakly backreacting nor under loop control. Large g_h does not by itself invalidate the exact solution of a quadratic mode equation on a prescribed background. The fixed-background approximation fails independently because of the energy

overdraw. Reducing the free quartic can cure both tree-level diagnostics along this parameter slice, but it changes the absolute scale $\mu \propto \lambda^{1/4}$ at fixed vacuum density and requires new Hubble, production, and loop checks.

5.7 Self-consistent semiclassical system

A minimal next calculation must solve the background, modes, and geometry together. At finite depth, one suitable Hartree system is

$$\ddot{\phi} + 3H\dot{\phi} + U'(\phi) + \frac{1}{2}F'(\phi)Y_H^{\mathrm{ren}} = 0, \tag{5.27}$$

$$\ddot{u}_{h,k} + 3H\dot{u}_{h,k} + \left(\frac{k^2}{a^2} + m_b^2 + F(\phi)\lambda_h\right) u_{h,k} = 0, \tag{5.28}$$

$$3M_{\mathrm{Pl}}^2 H^2 = \frac{1}{2}\dot{\phi}^2 + U(\phi) + \rho_H^{\mathrm{ren}} + \rho_{\mathrm{vis}}, \tag{5.29}$$

where

$$Y_H^{\mathrm{ren}} = \sum_{h=1}^{H} d_h \lambda_h \int \frac{\mathrm{d}^3k}{(2\pi)^3} \left(|u_{h,k}|^2 - \text{adiabatic subtraction}\right). \tag{5.30}$$

The same subtraction prescription must be used in ρ_H^{ren}. Equations (2.24)–(2.25) then provide an exact numerical conservation test.

The calculation is not complete until it demonstrates all of the following.

1. cutoff stability over a declared range of H and momentum resolution
2. conservation after renormalization
3. independence from an arbitrary scale-factor normalization
4. a spectral output rather than a shape-preserving projection
5. a regime in which quantum and higher-order corrections are controlled

No result in this chapter claims that such a solution exists for the frozen parameters. The energy and interaction-strength diagnostics make survival of that exact benchmark unlikely without a substantial microscopic revision.

CHAPTER 6

Finite-Depth Quantum Stability and Ultraviolet Growth

6.1 Why finite depth is not optional

The classical fiber operator is well defined on suitable infinite-depth domains. A quantum effective potential asks a different question. Every level contributes vacuum fluctuations, and the multiplicity grows with depth. Normal ordering a produced-particle density does not remove the vacuum terms in the effective action.

The honest starting point is therefore finite H with a declared physical or effective cutoff. Let

$$M_h^2(\phi) = m_b^2 + F(\phi)\lambda_h, \qquad \lambda_h = \Lambda^2 p^{2\sigma h}. \tag{6.1}$$

6.2 One-loop effective potential

For canonical real bosonic detail fields on a constant background with $M_h^2(\phi) > 0$, the flat-space, local-potential contribution in a standard subtraction scheme is

$$\boxed{\Delta U_H(\phi) = \frac{1}{64\pi^2}\sum_{h=1}^{H} d_h M_h^4(\phi)\left[\log\frac{M_h^2(\phi)}{\bar{\mu}^2} - \frac{3}{2}\right].} \tag{6.2}$$

At the locked origin, $F'(0) = 0$ and $F''(0) = 2\gamma/v^2$. The induced average curvature is

$$\boxed{\Delta m_\phi^2 = \frac{\gamma}{16\pi^2 v^2}\sum_{h=1}^{H} d_h \lambda_h M_{h0}^2\left[\log\frac{M_{h0}^2}{\bar{\mu}^2} - 1\right],} \tag{6.3}$$

where $M_{h0}^2 = m_b^2 + \lambda_h$.

Equation (6.3) does not by itself prove that locking fails. Counterterms can be chosen at finite depth. It proves that the locking curvature is radiatively sensitive to the tower unless a symmetry or ultraviolet completion controls the sum. On a curved background, additional $M_h^2 R$ and curvature-squared counterterms are required. If a level mass becomes negative, the logarithm develops an imaginary part and the expression is not a real stability potential.

6.3 Infinite-depth growth theorem

Theorem 6.1 (Bosonic tower ultraviolet growth). *Fix ϕ with $F(\phi) > 0$. For $\sigma > 0$ and large depth,*

$$\sum_{h=1}^{H} d_h M_h^4 \asymp F(\phi)^2 \Lambda^4 p^{(1+4\sigma)H}. \tag{6.4}$$

Writing

$$M_{\mathrm{UV}} = \Lambda p^{\sigma H}, \tag{6.5}$$

the leading power is

$$\boxed{\Delta U_H \sim \frac{M_{\mathrm{UV}}^{4+1/\sigma}}{\Lambda^{1/\sigma}}} \tag{6.6}$$

up to logarithms and scheme-dependent coefficients.

Proof. At high level, $d_h \sim p^h$ and $M_h^4 \sim F(\phi)^2 \Lambda^4 p^{4\sigma h}$. Their product is a geometric sequence with ratio $p^{1+4\sigma}$. The last term controls the finite sum. Substituting $p^H = (M_{\mathrm{UV}}/\Lambda)^{1/\sigma}$ gives (6.6), with the fixed factor $F(\phi)^2$ suppressed there. If $F = 0$, the base mass instead gives growth proportional to p^H. If $F < 0$, sufficiently high levels are tachyonic. □

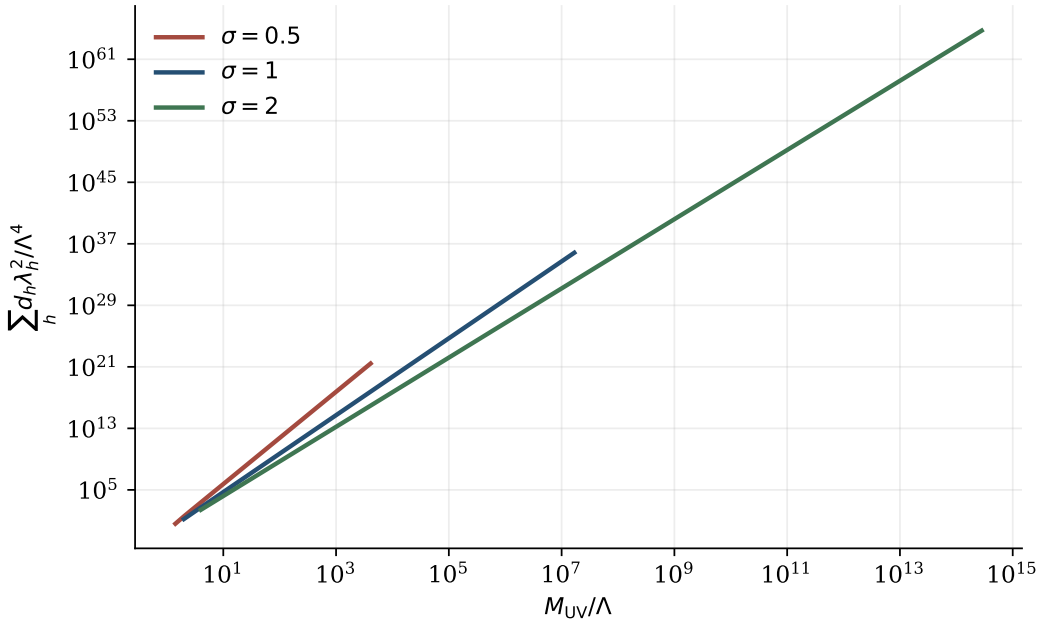


Figure 6.1. The multiplicity-weighted fourth mass moment at finite depth. The fitted slopes agree with $4 + 1/\sigma$. The lines have different horizontal ranges because the same depth produces different ultraviolet masses.

Ultraviolet boundary

The ordinary bosonic infinite tower does not have a finite one-loop potential without a physical cutoff, cancellations, or a different spectral ultraviolet prescription. Classical Hilbert-space convergence does not supply that cancellation.

6.4 Radiative closure of the one-coupling pattern

At tree level, the level interaction coefficients obey

$$g_h = \gamma \frac{\lambda_h}{v^2}. \tag{6.7}$$

A generic finite multiscalar quantum theory produces counterterms of the forms

$$\phi^4, \qquad \chi_h^4, \qquad \chi_h^2 \chi_j^2, \qquad \phi^2 \chi_h^2, \tag{6.8}$$

with independent finite parts unless a symmetry relates them. Level-orthogonal invariance constrains indices within one level. Branch homogeneity relates levels only if the regulator and renormalization prescription preserve the branching maps. A finite upper boundary introduces additional endpoint terms.

The single γ statement applies only to the order-2σ spectral channel of the classical ansatz. The allowed order-zero coefficient κ was set to zero as a model restriction, and generic loop corrections need not preserve that restriction. There is no renormalization theorem yet. A completed model needs one of the following:

1. a physical finite-depth endpoint
2. paired degrees of freedom that cancel the leading vacuum moments
3. an ultraviolet fixed point with a proved invariant counterterm family
4. a modulus or gauge principle that turns the coefficient pattern into a protected geometric quantity

6.5 A dynamical refinement-scale modulus

One candidate completion promotes the multiplier of L_H to a field. Introduce a canonical scalar s and define

$$\begin{aligned} S_H^{\mathrm{mod}} = \int \sqrt{-g}\, \mathrm{d}^4x \Bigg[& \frac{M_{\mathrm{Pl}}^2}{2} R - \frac{1}{2}(\nabla\phi)^2 - U(\phi) - \frac{1}{2}(\nabla s)^2 - W(s,\phi) \\ & - \frac{1}{2}\langle \nabla\chi, \nabla\chi \rangle_H - \frac{1}{2}\left\langle \chi, \left(m_b^2 + \mathrm{e}^{2s/f_s} L_H\right)\chi \right\rangle_H \Bigg]. \end{aligned} \tag{6.9}$$

Choose

$$W(s,\phi) = \frac{M_s^2}{2}\left[s - \frac{f_s}{2}\log F(\phi) \right]^2. \tag{6.10}$$

Proposition 6.2 (Heavy-modulus matching)**.** *If*

$$M_s^2 \gg H^2, \qquad M_s^2 \gg A_\star m_{\mathrm{peak}}, \qquad F > 0, \qquad |\Box s_0| \ll f_s M_s^2, \qquad F Y_H \ll f_s^2 M_s^2, \tag{6.11}$$

where $s_0(\phi) = \frac{f_s}{2}\log F(\phi)$*. Then the modulus equation has the adiabatic solution*

$$s = s_0 + \frac{\Box s_0}{M_s^2} - \frac{F Y_H}{f_s M_s^2} + \mathcal{O}(M_s^{-4}), \tag{6.12}$$

and hence

$$e^{2s/f_s} = F(\phi)\left[1 + \mathcal{O}\left(\frac{\Box s_0}{f_s M_s^2}, \frac{FY_H}{f_s^2 M_s^2}\right)\right]. \tag{6.13}$$

Proof. Expand the modulus equation around the moving minimum of (6.10). To leading order it is $\Box s_0 - M_s^2(s - s_0) - FY_H/f_s = 0$, which gives (6.12). Expanding the exponential gives (6.13). □

Model status

The proposition matches the detail mass multiplier, not the complete original scalar EFT. Substituting the moving minimum also produces the unsuppressed operator

$$-\frac{1}{2}(\nabla s_0)^2 = -\frac{f_s^2}{8}(\nabla \log F)^2. \tag{6.14}$$

It must be retained or shown negligible. The chosen potential (6.10) remains a model input. It does not derive γ, and it does not cancel the tower vacuum energy by itself.

6.6 Quantum tasks that remain

The next quantum block must supply a regulator compatible with diffeomorphism and internal symmetries, a finite counterterm basis, a measure contribution in the BV pushforward, and a cutoff-stable prediction. It must also decide whether H is physical, emergent, or removed through a cancellation. Until then, infinite-depth statements should be restricted to the classical operator domains for which they were proved.

CHAPTER 7

Physical Interpretation and Unresolved Comparisons

7.1 Mathematical construction and physical validation

A coherent mathematical construction, a controlled physical model, and an empirically supported theory are different kinds of result. The finite-depth covariant and canonical identities concern a stated model. The production spectrum additionally requires a controlled solution of the coupled dynamics. Empirical identification requires comparison with data and competing models.

The energy-budget obstruction locates a failure in the specified production history. It does not invalidate projector reconstruction, matrix-corner calculus, or the conditional transport of a prescribed distribution. It does prevent that distribution from being called a prediction of the microscopic crossing without a new dynamical argument.

7.2 Established model statements

The central results separate symmetry, conservation, approximation, and physical production.

Result	Status	Scientific role
Level-orthogonal quadratic operator classification	theorem	isolates one coefficient per refinement level
Branch-homogeneity recurrence	theorem	reduces those coefficients to one geometric family
Leading analytic spectral interaction	theorem within stated parity, locality, and exact-order assumptions	explains one channel but leaves an allowed order-zero coupling
Covariant stress-exchange identity	theorem	supplies an exact conservation diagnostic
Total hypersurface algebra with scalar fiber	theorem	shows that the finite-depth matter sector is canonically compatible after gravity is included
Resolved matter H–H test and negative controls	controlled diagnostic	separates spectral headroom from aliasing and omitted-gradient failures
Arbitrary-order continuum moments	theorem for the interpolated density model	extends the production ratios beyond $\sigma = 1$
Fixed-prime Mellin decomposition	theorem	corrects the interpretation of the continuum approximation
Specified benchmark energy audit	controlled diagnostic	proves the prescribed crossing is not weakly backreacting

Finite-depth one-loop power counting	theorem at leading ultraviolet order	exposes the need for a cutoff or cancellation
Heavy-modulus realization	model proposition	supplies a testable route to a common scale multiplier

The energy-budget result is a central obstruction. The specified rapid-crossing spectrum is not a small perturbation of its background. Its unprojected energy exceeds the declared reservoir by $2.6272875374 \times 10^{20}$. A free crossing followed by a common normalization factor therefore cannot be read as a self-consistent production history.

This negative result is constructive. It localizes the failure upstream of the late transport calculation and gives exact equations for replacing the projection. It also prevents further phenomenology from accumulating on an uncontrolled microscopic input.

7.3 Limits of the model

The following claims remain outside the proved domain.

1. The projector axioms do not select p, σ, Λ, m_b, γ, the potential parameters, or the initial state.
2. The symmetry classification does not calculate the numerical value $\gamma = 16$.
3. The order-one backreaction threshold is not a solution of the coupled background and mode equations or an energy-partition equation.
4. The projected Gaussian ratios are not preserved under strong backreaction unless a dynamical calculation demonstrates that fact.
5. The ordinary infinite bosonic tower is not radiatively finite under the finite-depth continuation studied here.
6. The modulus completion does not solve the cosmological constant problem or prove ultraviolet completeness.
7. The internal verification archive is not independent reproduction, peer review, or experimental evidence.
8. No dark matter or dark energy identification follows solely from a match to a prescribed background abundance.

7.4 A controlled production comparison

An additional observable template does not resolve the upstream production problem. It is a finite-depth backreacting production calculation with one conserved renormalized stress tensor. The system in (5.27)–(5.29) is the minimum viable model.

The calculation should be organized as a sequence of gates.

Gate	Required result	Failure meaning

G1	Renormalized total continuity residual converges to zero with time step and momentum resolution	implementation or subtraction is inconsistent
G2	Observables stabilize across at least three finite depths and two momentum cutoffs	prediction is regulator dominated
G3	The produced spectrum is obtained from the coupled evolution without a common post hoc projection	microscopic production is actually solved
G4	Interaction and loop expansion parameters are below declared bounds, or a nonperturbative method is used	free-mode interpretation is invalid
G5	A viable region reproduces abundance, pressure, and coldness simultaneously	the proposed cosmological branch survives
G6	At least one linked observable differs from a simpler model after nuisance parameters are fit	the construction becomes falsifiable rather than only flexible

Go or revise

If no cutoff-stable solution passes G1–G5, the specified cosmological branch should be revised or retired. The projector, corner, and constrained-system results would remain as separate mathematical work. A failed dark-sector realization would not invalidate them.

7.5 Minimum computational experiment

A credible first run does not need the full infinite tower. It needs a converged finite-depth sequence.

1. Choose $H = 1, 2, 3, 4$ and publish the physical largest mass at every depth.
2. Initialize fourth-order adiabatic vacuum modes over a documented comoving momentum grid.
3. Use the same adiabatic subtraction terms in Y_H^{ren}, ρ_H^{ren}, and p_H^{ren}.
4. Evolve ϕ, all retained modes, and the Friedmann equation together through the crossing.
5. Monitor the covariant identity in (2.19) and the FLRW exchanges in (2.20)–(2.21) at every step.
6. Repeat with doubled time resolution, doubled momentum resolution, and shifted ultraviolet cutoffs.
7. Report the unprojected spectrum, energy, pressure, rest fraction, level occupancy, and their convergence envelopes.
8. Scan λ, γ, Λ/μ, m_b/μ, and the invariant crossing condition. Do not use a coordinate value of the scale factor as the only clock.

The first scientific deliverable from this experiment should be a survival map, not a best-fit point. Its axes should distinguish the weakly coupled, strongly backreacting, tachyonic, cutoff-sensitive, and radiatively unstable regions.

7.6 A second foundational track

The strongest purely gravitational continuation is different. It is an end-to-end nonlinear Gowdy corner–BFV continuum bridge. That track should derive the ungauge-fixed T^2-symmetric sector from the real polynomial Lorentzian parent, construct an explicit corner derivative and variable-shift generator, instantiate a cyclic contraction, prove finite BFV nilpotency and refinement-path equivalence, and establish uniform local-in-time convergence after areal gauge on $[t_0, T]$ with $t_0 > 0$.

That theorem would close the gap between the abstract matrix-corner calculus and the filtered Fourier realizations used in the later gravitational chapters. It is the right foundational milestone alongside the matter comparison. It should not be replaced by an unrestricted $3 + 1$ or quantum claim before the symmetry-reduced continuum map is explicit.

7.7 Decision table for the program

Outcome	Interpretation	Next action
Backreacting spectrum remains close to the prescribed ansatz in a controlled region	cosmological bridge gains microscopic support	rerun late observables with propagated theory uncertainty
Spectrum survives but is substantially reshaped	mechanism may survive, benchmark does not	replace frozen initial distribution and refit every downstream stage
Only strongly coupled solutions exist	free production calculation is the wrong tool	formulate and validate a nonperturbative treatment
Finite-depth results fail to stabilize	tower prediction is regulator dependent	identify a physical endpoint or ultraviolet cancellation
No viable production region exists	dark-sector realization is falsified in this form	retain mathematical framework and revise the matter completion
Gowdy bridge succeeds independently	foundational architecture gains a nonlinear continuum theorem	extend cautiously to additional symmetry classes

7.8 Final claim boundary

The construction supplies exact internal identities, conditional gravitational statements, and an explicitly testable fiber model. Its specified cosmological production benchmark fails a quantified energy-budget test. A viable physical realization must control backreaction, ultraviolet sensitivity, and the transport of any resulting uncertainty through the observable calculations.

The cosmology chapter therefore begins from a prescribed normalized distribution. Its deductions retain their conditional meaning. They do not assume that the unresolved production comparison has succeeded.

Cosmology, Dark Sectors, and Structure Formation

Conditional transport from a prescribed refinement distribution

The preceding fiber model separates a candidate covariant action from the production history used to populate it. This distinction controls the interpretation of the cosmological calculations. The specified rapid-crossing benchmark fails its available-energy test. A common post hoc normalization is not a solution of the coupled background and mode equations.

> **The distribution is an input**
>
> The transport, CMB, structure, baryonic, and subhalo calculations below are conditional on a prescribed normalized distribution and on their stated approximations. They are not predictions of a demonstrated microscopic production process. The exact field identities and bounded calculations retain their mathematical meaning under that condition.

The account first fixes relativistic background and perturbation conventions. It then defines the refinement sector, propagates its specified distribution through kinetic and observable models, and identifies linked rejection tests. Numerical controls apply to the declared implementations and parameter choices. They do not establish precision agreement with data or empirical preference over simpler models.

Symbols and Conventions

Symbol	Meaning
$a, H, \mathcal{H}$	Scale factor, physical Hubble rate, and conformal Hubble rate.
Φ, Ψ	Scalar metric potentials.
ϕ	Fiber-average field.
$\chi_{h,a}$	Haar-detail field.
D_h	Haar-detail projector.
Λ, σ, p	Refinement mass scale, order, and prime branching parameter.
γ	Scale-covariant coupling.
$a_\star, A_\star, \nu_h$	Production epoch, crossing scale, and level production scale.
ρ_D, p_D, w_D	Detail-tower density, pressure, and equation of state.
M_ϕ^2	Average-mode locking curvature.
$F_{h\ell}$	Collisionless multipoles.
C_ℓ^{XY}, D_ℓ^{XY}	Angular spectra and $\ell(\ell+1)C_\ell/(2\pi)$.
$T_D(k)$	Refinement-to-cold transfer amplitude.
$M_{1/2}, k_{1/2}$	Half-mode mass and wavenumber.

CHAPTER 1

Curved Multicomponent FLRW Background and Distances

Background geometry and perturbations. The first part fixes the common FLRW and perturbation conventions used later. It is broader than the dark-sector model and supplies the transport language in which model-specific effects are measured.

The homogeneous model is extended to spatial curvature, a cosmological constant, multiple fluids, and standard distance observables. The equations define a controlled background solver rather than a microscopic derivation of every component.

1.1 Curved lapse-resolved minisuperspace action

Let

$$\mathrm{d}s^2 = -N(t)^2\,\mathrm{d}t^2 + a(t)^2\gamma^{(K)}_{ij}\,\mathrm{d}x^i\,\mathrm{d}x^j, \qquad {}^{(3)}R[\gamma^{(K)}] = 6K. \tag{1.1}$$

For unit comoving volume,

$$S_{\mathrm{FLRW}} = \int \left[-\frac{3M_{\mathrm{Pl}}^2 a\dot{a}^2}{N} + 3M_{\mathrm{Pl}}^2 KaN - M_{\mathrm{Pl}}^2\Lambda a^3 N - NH_{\mathrm{P}}(a)\right]\mathrm{d}t. \tag{1.2}$$

Varying N gives

$$0 = \frac{3M_{\mathrm{Pl}}^2 a\dot{a}^2}{N^2} + 3M_{\mathrm{Pl}}^2 Ka - M_{\mathrm{Pl}}^2\Lambda a^3 - H_{\mathrm{P}}(a). \tag{1.3}$$

After division by $3M_{\mathrm{Pl}}^2 a^3$ and use of $H = \dot{a}/(Na)$, this becomes the Friedmann equation. With

$$\rho(a) := \frac{H_{\mathrm{P}}(a)}{a^3}, \tag{1.4}$$

the lapse equation gives

$$\boxed{H^2 = \frac{\rho}{3M_{\mathrm{Pl}}^2} + \frac{\Lambda}{3} - \frac{K}{a^2}.} \tag{1.5}$$

For total pressure p,

$$\dot{H} = -\frac{\rho + p}{2M_{\mathrm{Pl}}^2} + \frac{K}{a^2}, \tag{1.6}$$

$$\frac{\ddot{a}}{a} = -\frac{\rho + 3p}{6M_{\mathrm{Pl}}^2} + \frac{\Lambda}{3}, \tag{1.7}$$

$$\dot{\rho} + 3H(\rho + p) = 0. \tag{1.8}$$

In the gauge $N = 1$, differentiate the Friedmann equation:

$$2H\dot{H} = \frac{\dot{\rho}}{3M_{\mathrm{Pl}}^2} + \frac{2KH}{a^2}. \tag{1.9}$$

Using $\dot{\rho} = -3H(\rho + p)$ and dividing by $2H$ gives the equation for $\dot{H}$. Then

$$\frac{\ddot{a}}{a} = \dot{H} + H^2 \tag{1.10}$$

$$= -\frac{\rho + p}{2M_{\mathrm{Pl}}^2} + \frac{K}{a^2} + \frac{\rho}{3M_{\mathrm{Pl}}^2} + \frac{\Lambda}{3} - \frac{K}{a^2}, \tag{1.11}$$

which simplifies to the acceleration equation. The symbolic identity

$$\frac{a}{2}\frac{\mathrm{d}}{\mathrm{d}a}\left[\frac{8\pi G}{3}\rho + \frac{\Lambda}{3} - \frac{K}{a^2}\right] = -4\pi G(\rho + p) + \frac{K}{a^2} \tag{1.12}$$

is exact when

$$a\rho_{,a} = -3(\rho + p). \tag{1.13}$$

1.2 Constant-equation-of-state component library

For noninteracting components

$$p_i = w_i\rho_i, \tag{1.14}$$

$$\dot{\rho}_i + 3H(1 + w_i)\rho_i = 0, \tag{1.15}$$

$$\rho_i(a) = \rho_{i0}a^{-3(1+w_i)}. \tag{1.16}$$

Since $\dot{\rho}_i = aH\ \mathrm{d}\rho_i/\ \mathrm{d}a$,

$$\frac{\mathrm{d}\rho_i}{\rho_i} = -3(1 + w_i)\frac{\mathrm{d}a}{a}. \tag{1.17}$$

Integration from $a = 1$ to a gives the stated power law. With $a_0 = 1$,

$$\Omega_{i0} := \frac{\rho_{i0}}{3M_{\mathrm{Pl}}^2H_0^2}, \tag{1.18}$$

$$\Omega_{\Lambda 0} := \frac{\Lambda}{3H_0^2}, \tag{1.19}$$

$$\Omega_{K0} := -\frac{K}{H_0^2}, \tag{1.20}$$

$$1 = \sum_i \Omega_{i0} + \Omega_{K0} + \Omega_{\Lambda 0}, \tag{1.21}$$

$$E(a)^2 := \frac{H(a)^2}{H_0^2} = \sum_i \Omega_{i0}a^{-3(1+w_i)} + \Omega_{K0}a^{-2} + \Omega_{\Lambda 0}. \tag{1.22}$$

Hence

$$E(z)^2 = \sum_i \Omega_{i0}(1 + z)^{3(1+w_i)} + \Omega_{K0}(1 + z)^2 + \Omega_{\Lambda 0}. \tag{1.23}$$

Special cases:

$$\rho_r \propto a^{-4}, \tag{1.24}$$
$$\rho_m \propto a^{-3}, \tag{1.25}$$
$$\rho_K \propto a^{-2}, \tag{1.26}$$
$$\rho_\Lambda = \text{constant}. \tag{1.27}$$

1.3 Conservative interacting fluids

Let

$$\dot{\rho}_i + 3H(\rho_i + p_i) = Q_i, \qquad \sum_i Q_i = 0. \tag{1.28}$$

Then

$$\boxed{\dot{\rho} + 3H(\rho + p) = 0, \qquad \rho := \sum_i \rho_i, \quad p := \sum_i p_i.} \tag{1.29}$$

Summing the component equations gives

$$\dot{\rho} + 3H(\rho + p) = \sum_i Q_i = 0. \tag{1.30}$$

Energy transfer therefore changes the partition among components but not the total homogeneous continuity equation. For the benchmark transfer

$$Q_m = -\Gamma\rho_m, \tag{1.31}$$
$$Q_r = +\Gamma\rho_m, \tag{1.32}$$

$$\dot{\rho}_m = -3H\rho_m - \Gamma\rho_m, \tag{1.33}$$
$$\dot{\rho}_r = -4H\rho_r + \Gamma\rho_m, \tag{1.34}$$
$$\dot{\rho}_m + \dot{\rho}_r + 3H\left(\rho_m + \rho_r + \frac{1}{3}\rho_r\right) = 0. \tag{1.35}$$

1.4 Background observables

Define

$$\bar{\chi}(z) := \int_0^z \frac{\mathrm{d}z'}{E(z')}, \tag{1.36}$$
$$D_H := \frac{c}{H_0}. \tag{1.37}$$

For a radial null ray,

$$0 = -c^2\,\mathrm{d}t^2 + a(t)^2\,\mathrm{d}\chi_{\mathrm{com}}^2. \tag{1.38}$$

Using $1 + z = a^{-1}$ and $\mathrm{d}t = -\,\mathrm{d}z/[(1+z)H(z)]$ gives

$$\mathrm{d}\chi_{\mathrm{com}} = \frac{c\,\mathrm{d}z}{H(z)} = D_H\frac{\mathrm{d}z}{E(z)}. \tag{1.39}$$

Thus $D_H\bar{\chi}(z)$ is the radial comoving distance before the curvature map is applied. The transverse comoving distance is

$$D_M(z) = D_H\,\mathcal{S}_{\Omega_K}(\bar{\chi}(z))\,, \tag{1.40}$$

where

$$\mathcal{S}_{\Omega_K}(x) := \begin{cases} \Omega_K^{-1/2}\sinh(\sqrt{\Omega_K}\,x), & \Omega_K > 0, \\ x, & \Omega_K = 0, \\ |\Omega_K|^{-1/2}\sin(\sqrt{|\Omega_K|}\,x), & \Omega_K < 0. \end{cases} \tag{1.41}$$

Then

$$D_A(z) = \frac{D_M(z)}{1+z}, \tag{1.42}$$

$$D_L(z) = (1+z)D_M(z), \tag{1.43}$$

$$\boxed{D_L(z) = (1+z)^2 D_A(z)}, \tag{1.44}$$

$$H_0 t_{\mathrm{L}}(z) = \int_0^z \frac{\mathrm{d}z'}{(1+z')E(z')}, \tag{1.45}$$

$$H_0 t_0 = \int_0^1 \frac{\mathrm{d}a}{aE(a)}. \tag{1.46}$$

Exact benchmarks follow by direct integration. For matter domination, $E(a) = a^{-3/2}$ and $E(z) = (1+z)^{3/2}$, so

$$H_0 t_0 = \int_0^1 a^{1/2}\,\mathrm{d}a = \frac{2}{3}, \tag{1.47}$$

$$\bar{\chi}(z) = \int_0^z (1+z')^{-3/2}\,\mathrm{d}z' \tag{1.48}$$

$$= 2\left[1 - (1+z)^{-1/2}\right]. \tag{1.49}$$

For radiation domination, $E(a) = a^{-2}$ and $E(z) = (1+z)^2$, giving

$$H_0 t_0 = \int_0^1 a\;\mathrm{d}a = \frac{1}{2}, \tag{1.50}$$

$$\bar{\chi}(z) = \int_0^z \frac{\mathrm{d}z'}{(1+z')^2} = \frac{z}{1+z}. \tag{1.51}$$

For pure curvature, $E(a) = a^{-1}$ and $E(z) = 1+z$, so

$$H_0 t_0 = \int_0^1 \mathrm{d}a = 1, \tag{1.52}$$

$$\bar{\chi}(z) = \int_0^z \frac{\mathrm{d}z'}{1+z'} = \log(1+z), \tag{1.53}$$

$$\frac{D_M}{D_H} = \sinh\bar{\chi} \qquad (\Omega_K = 1). \tag{1.54}$$

1.5 Projector background solver

With normalized time $\tau = H_0 t$,

$$\frac{\mathrm{d}a}{\mathrm{d}\tau} = aE(a). \tag{1.55}$$

The explicit midpoint update is

$$k_1 = a_n E(a_n), \tag{1.56}$$

$$k_2 = \left(a_n + \frac{\Delta\tau}{2}k_1\right) E\left(a_n + \frac{\Delta\tau}{2}k_1\right), \tag{1.57}$$

$$a_{n+1} = a_n + \Delta\tau k_2. \tag{1.58}$$

For the curved radiation–matter–Λ benchmark,

$$p_a = 2.024278. \tag{1.59}$$

For conservative matter-to-radiation transfer,

$$p_{\text{transfer}} = 2.054690. \tag{1.60}$$

The recorded residuals are

$$\varepsilon_{\text{distance benchmarks}} = 8.881784 \times 10^{-16}, \tag{1.61}$$

$$\varepsilon_{\text{age benchmarks}} = 0, \tag{1.62}$$

$$\varepsilon_{D_L-(1+z)^2 D_A} = 0, \tag{1.63}$$

$$\varepsilon_{\text{total interacting continuity}} = 4.440892 \times 10^{-16}. \tag{1.64}$$

1.6 Open extensions

- Derive a microscopic action that produces the exchange terms Q_i rather than inserting them phenomenologically.
- Prove positivity, regularity, and global-existence conditions for the interacting background system.
- Extend the projector background equations to anisotropic Bianchi geometries.
- Determine whether nonsingular bouncing solutions exist and establish their dynamical stability when they do.
- Calibrate the resulting background model against cosmological distance and expansion data.

CHAPTER 2

Gauge-Invariant Einstein–Scalar Perturbations

Scalar perturbations are reduced to a single gauge-invariant degree of freedom. The continuum Mukhanov–Sasaki action is then discretized with the same projector Laplacian used elsewhere.

2.1 Einstein–scalar input and background

Let

$$M_{\mathrm{Pl}}^2 := (8\pi G)^{-1} \tag{2.1}$$

and

$$S[g,\phi] = \int \sqrt{-g}\left[\frac{M_{\mathrm{Pl}}^2}{2}R - \frac{1}{2}g^{\mu\nu}\partial_\mu\phi\partial_\nu\phi - V(\phi)\right]\mathrm{d}^4x. \tag{2.2}$$

For

$$\mathrm{d}s^2 = -\,\mathrm{d}t^2 + a(t)^2\,\mathrm{d}\boldsymbol{x}^2, \qquad \phi = \bar{\phi}(t), \tag{2.3}$$

the background equations are

$$3M_{\mathrm{Pl}}^2H^2 = \frac{1}{2}\dot{\bar{\phi}}^2 + V(\bar{\phi}), \tag{2.4}$$

$$-2M_{\mathrm{Pl}}^2\dot{H} = \dot{\bar{\phi}}^2, \tag{2.5}$$

$$\ddot{\bar{\phi}} + 3H\dot{\bar{\phi}} + V_{,\phi} = 0, \tag{2.6}$$

$$\epsilon := -\frac{\dot{H}}{H^2} = \frac{\dot{\bar{\phi}}^2}{2M_{\mathrm{Pl}}^2H^2}. \tag{2.7}$$

2.2 Scalar gauge invariant

Use the scalar convention

$$g_{ij} = a^2\left[(1-2\Psi)\delta_{ij} + 2\partial_i\partial_jE\right]. \tag{2.8}$$

Under $t \mapsto t+T$,

$$\Psi \mapsto \Psi + HT, \tag{2.9}$$

$$\delta\phi \mapsto \delta\phi - \dot{\bar{\phi}}T. \tag{2.10}$$

Define

$$\boxed{\zeta := -\Psi - \frac{H}{\dot{\bar{\phi}}}\delta\phi.} \tag{2.11}$$

Indeed, substitution of the two transformations gives

$$\zeta \mapsto -(\Psi + HT) - \frac{H}{\dot{\phi}}\left(\delta\phi - \dot{\phi}T\right) \tag{2.12}$$

$$= -\Psi - HT - \frac{H}{\dot{\phi}}\delta\phi + HT \tag{2.13}$$

$$= \zeta. \tag{2.14}$$

The time-shift contributions therefore cancel before any gauge choice is made. In comoving gauge,

$$\delta\phi = 0, \qquad E = 0, \qquad h_{ij} = a^2 \mathrm{e}^{2\zeta}\delta_{ij}. \tag{2.15}$$

2.3 Lapse and shift elimination

Write

$$N = 1 + \alpha, \qquad N_i = \partial_i \beta. \tag{2.16}$$

The linear Hamiltonian and momentum constraints give

$$\alpha = \frac{\dot{\zeta}}{H}, \tag{2.17}$$

$$\beta = -\frac{\zeta}{H} + \chi, \tag{2.18}$$

$$\nabla^2 \chi = a^2 \epsilon \dot{\zeta}. \tag{2.19}$$

Thus α and β are nondynamical and are eliminated before the reduced quadratic action is formed.

2.4 Reduced scalar action and equation

After substitution of the constraints,

$$\boxed{S_\zeta^{(2)} = M_{\mathrm{Pl}}^2 \int a^3 \epsilon \left[\dot{\zeta}^2 - \frac{(\nabla\zeta)^2}{a^2}\right] \mathrm{d}t\, \mathrm{d}^3x.} \tag{2.20}$$

For a compactly supported variation $\delta\zeta$,

$$\delta S_\zeta^{(2)} = 2M_{\mathrm{Pl}}^2 \int \left[a^3 \epsilon \dot{\zeta}\, \delta\dot{\zeta} - a\epsilon \partial_i\zeta\, \partial_i\delta\zeta\right] \mathrm{d}t\, \mathrm{d}^3x \tag{2.21}$$

$$= -2M_{\mathrm{Pl}}^2 \int \left[\frac{\mathrm{d}}{\mathrm{d}t}\left(a^3\epsilon\dot{\zeta}\right) - a\epsilon\nabla^2\zeta\right] \delta\zeta\, \mathrm{d}t\, \mathrm{d}^3x, \tag{2.22}$$

where the second line follows by integrating once in time and once in space. Hence the Euler–Lagrange equation is

$$\boxed{\frac{\mathrm{d}}{\mathrm{d}t}\left(a^3\epsilon\dot{\zeta}\right) - a\epsilon\nabla^2\zeta = 0.} \tag{2.23}$$

For a Fourier mode, $\nabla^2\zeta_{\boldsymbol{k}} = -k^2\zeta_{\boldsymbol{k}}$, so

$$0 = \frac{\mathrm{d}}{\mathrm{d}t}\left(a^3\epsilon\dot{\zeta}_{\boldsymbol{k}}\right) + a\epsilon k^2\zeta_{\boldsymbol{k}} \tag{2.24}$$

$$= a^3\epsilon\ddot{\zeta}_{\boldsymbol{k}} + \left(3Ha^3\epsilon + a^3\dot{\epsilon}\right)\dot{\zeta}_{\boldsymbol{k}} + a\epsilon k^2\zeta_{\boldsymbol{k}}. \tag{2.25}$$

Division by $a^3\epsilon$ gives

$$\ddot{\zeta}_{\boldsymbol{k}} + \left(3H + \frac{\dot{\epsilon}}{\epsilon}\right)\dot{\zeta}_{\boldsymbol{k}} + \frac{k^2}{a^2}\zeta_{\boldsymbol{k}} = 0. \tag{2.26}$$

At $k = 0$, the equation can be integrated twice:

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(a^3\epsilon\dot{\zeta}\right) = 0, \tag{2.27}$$

$$a^3\epsilon\dot{\zeta} = C_2, \tag{2.28}$$

$$\dot{\zeta} = \frac{C_2}{a^3\epsilon}, \tag{2.29}$$

$$\zeta(t) = C_1 + C_2\int^t \frac{\mathrm{d}s}{a(s)^3\epsilon(s)}. \tag{2.30}$$

Hence the adiabatic constant mode is conserved when the second mode decays.

2.5 Mukhanov–Sasaki form

Let

$$\mathrm{d}t = a\,\mathrm{d}\eta, \tag{2.31}$$

$$z(\eta) := a\sqrt{2\epsilon}\,M_{\mathrm{Pl}}, \tag{2.32}$$

$$v := z\zeta. \tag{2.33}$$

Changing from t to η first gives

$$S_\zeta^{(2)} = \frac{1}{2}\int z^2\left[(\zeta')^2 - (\nabla\zeta)^2\right]\,\mathrm{d}\eta\,\mathrm{d}^3x. \tag{2.34}$$

Since $\zeta = v/z$ and z is spatially homogeneous,

$$z^2(\zeta')^2 = \left(v' - \frac{z'}{z}v\right)^2 \tag{2.35}$$

$$= (v')^2 - \frac{z'}{z}(v^2)' + \left(\frac{z'}{z}\right)^2 v^2, \tag{2.36}$$

$$z^2(\nabla\zeta)^2 = (\nabla v)^2. \tag{2.37}$$

Integrating the mixed term by parts yields

$$-\int \frac{z'}{z}(v^2)'\,\mathrm{d}\eta = \int\left(\frac{z'}{z}\right)' v^2\,\mathrm{d}\eta + \text{boundary term}, \tag{2.38}$$

$$\left(\frac{z'}{z}\right)' + \left(\frac{z'}{z}\right)^2 = \frac{z''}{z}. \tag{2.39}$$

Therefore, up to the displayed boundary term,

$$\boxed{S_v^{(2)} = \frac{1}{2}\int \left[(v')^2 - (\nabla v)^2 + \frac{z''}{z}v^2\right] \mathrm{d}\eta\, \mathrm{d}^3x.} \tag{2.40}$$

Therefore

$$v'' - \nabla^2 v - \frac{z''}{z}v = 0, \tag{2.41}$$

$$v_{\boldsymbol{k}}'' + \left(k^2 - \frac{z''}{z}\right) v_{\boldsymbol{k}} = 0. \tag{2.42}$$

2.6 Projector-lattice scalar action

For cubic spacing h,

$$S_{v,h}^{(2)} = \frac{1}{2}\int h^3 \sum_j \left[(v_j')^2 - \sum_{r=1}^{3}(\delta_r^+ v_j)^2 + \frac{z''}{z}v_j^2\right] \mathrm{d}\eta. \tag{2.43}$$

The discrete summation-by-parts identity

$$h^3 \sum_{j,r}(\delta_r^+ v_j)(\delta_r^+ \delta v_j) = -h^3 \sum_j (\Delta_h v_j)\delta v_j \tag{2.44}$$

converts the gradient variation into the lattice Laplacian. The lattice equation is therefore

$$v_j'' - \Delta_h v_j - \frac{z''}{z}v_j = 0. \tag{2.45}$$

For a lattice Fourier mode,

$$\delta_r^+ \mathrm{e}^{\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}_j} = \frac{\mathrm{e}^{\mathrm{i}k_r h} - 1}{h}\mathrm{e}^{\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}_j}, \tag{2.46}$$

$$\left|\frac{\mathrm{e}^{\mathrm{i}k_r h} - 1}{h}\right|^2 = \frac{(\mathrm{e}^{\mathrm{i}k_r h} - 1)(\mathrm{e}^{-\mathrm{i}k_r h} - 1)}{h^2} \tag{2.47}$$

$$= \frac{2 - 2\cos(k_r h)}{h^2} = \frac{4}{h^2}\sin^2\left(\frac{k_r h}{2}\right). \tag{2.48}$$

Thus

$$\widehat{k}_h^2 := \frac{4}{h^2}\sum_{r=1}^{3}\sin^2\left(\frac{k_r h}{2}\right), \tag{2.49}$$

$$v_{\boldsymbol{k},h}'' + \left(\widehat{k}_h^2 - \frac{z''}{z}\right) v_{\boldsymbol{k},h} = 0, \tag{2.50}$$

$$\widehat{k}_h^2 = k^2 - \frac{h^2}{12}\sum_{r=1}^{3}k_r^4 + \frac{h^4}{360}\sum_{r=1}^{3}k_r^6 + \mathcal{O}(h^6). \tag{2.51}$$

For $z''/z = 2/\eta^2$,

$$v_\kappa(\eta) = \frac{1}{\sqrt{2\kappa}}\left(1 - \frac{\mathrm{i}}{\kappa\eta}\right)\mathrm{e}^{-\mathrm{i}\kappa\eta} \tag{2.52}$$

solves

$$v_\kappa'' + \left(\kappa^2 - \frac{2}{\eta^2}\right)v_\kappa = 0. \tag{2.53}$$

With $\kappa = \widehat{k}_h$,

$$p_{v,h} = 2.001500, \tag{2.54}$$

$$p_{|v|^2,h} = 2.002803, \tag{2.55}$$

$$p_{v,h}^{\mathrm{Richardson}} = 4.001170, \tag{2.56}$$

$$\max_h \left|v_h v_h^{*\prime} - v_h^* v_h' - \mathrm{i}\right| = 3.219647 \times 10^{-15}. \tag{2.57}$$

2.7 Open extensions

- Extend the quadratic derivation from one canonical scalar to general background and transfer functions.
- Add vector and tensor sources within the same gauge-invariant projector calculus.
- Generalize to multiple scalar fields and derive the coupled adiabatic and entropy modes.
- Derive the cubic action and the leading non-Gaussian correlation functions.
- Obtain the vacuum normalization from the projector dynamics rather than prescribing it as an external input.

CHAPTER 3

Vector and Tensor Perturbations on FLRW

The scalar–vector–tensor split separates decaying vector modes from propagating tensor modes on FLRW. The discrete tensor sector preserves the continuum causal structure in the refinement limit.

3.1 Scalar–vector–tensor split

In conformal time,

$$\mathrm{d}s^2 = a(\eta)^2 \left[-\mathrm{d}\eta^2 + 2S_i \,\mathrm{d}\eta \,\mathrm{d}x^i + \left(\delta_{ij} + 2\partial_{(i}F_{j)} + h_{ij} \right) \mathrm{d}x^i \,\mathrm{d}x^j \right], \tag{3.1}$$

with

$$\partial^i S_i = 0, \tag{3.2}$$
$$\partial^i F_i = 0, \tag{3.3}$$
$$\partial^i h_{ij} = 0, \tag{3.4}$$
$$h^i{}_i = 0. \tag{3.5}$$

Under a transverse spatial gauge shift L_i,

$$S_i \mapsto S_i - L_i', \tag{3.6}$$
$$F_i \mapsto F_i - L_i. \tag{3.7}$$

Therefore

$$\boxed{\Sigma_i := S_i - F_i'} \tag{3.8}$$

is gauge invariant, because

$$\Sigma_i \mapsto (S_i - L_i') - (F_i - L_i)' \tag{3.9}$$
$$= S_i - L_i' - F_i' + L_i' \tag{3.10}$$
$$= \Sigma_i. \tag{3.11}$$

3.2 Source-free vector sector

In the absence of vector anisotropic stress,

$$\boxed{(a^2\Sigma_i)' = 0.} \tag{3.12}$$

Integrating once gives

$$(a^2\Sigma_i)' = 0, \tag{3.13}$$

$$a^2\Sigma_i = C_i, \tag{3.14}$$

$$\Sigma_i(\eta) = \frac{C_i}{a(\eta)^2}. \tag{3.15}$$

For a matter background $a \propto \eta^2$,

$$\Sigma_i \propto \eta^{-4}. \tag{3.16}$$

A midpoint update of

$$\Sigma' = -\frac{4}{\eta}\Sigma \tag{3.17}$$

converges with

$$p_{\text{vector}} = 2.052951. \tag{3.18}$$

The fully matter-sourced vector constraint and vorticity hierarchy are not derived here.

3.3 Tensor action and equation

The quadratic tensor action is

$$\boxed{S_T^{(2)} = \frac{M_{\text{Pl}}^2}{8}\int a^2\left[(h_{ij}')^2 - (\partial_k h_{ij})^2\right]\,\mathrm{d}\eta\,\mathrm{d}^3x.} \tag{3.19}$$

For a transverse-traceless variation,

$$\delta S_T^{(2)} = \frac{M_{\text{Pl}}^2}{4}\int a^2\left[h_{ij}'\delta h_{ij}' - \partial_k h_{ij}\partial_k\delta h_{ij}\right]\,\mathrm{d}\eta\,\mathrm{d}^3x \tag{3.20}$$

$$= -\frac{M_{\text{Pl}}^2}{4}\int\left[(a^2h_{ij}')' - a^2\nabla^2 h_{ij}\right]\delta h_{ij}\,\mathrm{d}\eta\,\mathrm{d}^3x. \tag{3.21}$$

Since

$$(a^2h_{ij}')' = a^2h_{ij}'' + 2aa'h_{ij}' = a^2(h_{ij}'' + 2\mathcal{H}h_{ij}'), \tag{3.22}$$

the vacuum Euler–Lagrange equation is

$$\boxed{h_{ij}'' + 2\mathcal{H}h_{ij}' - \nabla^2 h_{ij} = 0, \qquad \mathcal{H} := a'/a.} \tag{3.23}$$

With transverse-traceless anisotropic stress,

$$h_{ij}'' + 2\mathcal{H}h_{ij}' - \nabla^2 h_{ij} = \frac{2a^2}{M_{\text{Pl}}^2}\Pi_{ij}^{\text{TT}}. \tag{3.24}$$

For each Fourier polarization, set

$$\mu_\lambda := \frac{aM_{\text{Pl}}}{2}h_\lambda, \qquad h_\lambda = \frac{2\mu_\lambda}{aM_{\text{Pl}}}. \tag{3.25}$$

Then

$$h_\lambda' = \frac{2}{aM_{\text{Pl}}}(\mu_\lambda' - \mathcal{H}\mu_\lambda), \tag{3.26}$$

$$h_\lambda'' = \frac{2}{aM_{\rm Pl}}\left[\mu_\lambda'' - 2\mathcal{H}\mu_\lambda' + (\mathcal{H}^2 - \mathcal{H}')\mu_\lambda\right]. \tag{3.27}$$

Substitution into $h_\lambda'' + 2\mathcal{H}h_\lambda' + k^2 h_\lambda = 0$ cancels the first-derivative terms and gives

$$0 = \mu_\lambda'' + \left[k^2 - (\mathcal{H}' + \mathcal{H}^2)\right]\mu_\lambda \tag{3.28}$$

$$= \mu_\lambda'' + \left(k^2 - \frac{a''}{a}\right)\mu_\lambda. \tag{3.29}$$

Thus

$$\boxed{\mu_\lambda'' + \left(k^2 - \frac{a''}{a}\right)\mu_\lambda = 0.} \tag{3.30}$$

3.4 Projector-lattice tensor sector

For cubic spacing h, the centered projector Laplacian has symbol

$$\widehat{k}_h^2 = \frac{4}{h^2}\sum_{r=1}^{3}\sin^2\left(\frac{k_r h}{2}\right) \tag{3.31}$$

$$= \sum_{r=1}^{3}\left[k_r^2 - \frac{h^2 k_r^4}{12} + \frac{h^4 k_r^6}{360} + \mathcal{O}(h^6)\right] \tag{3.32}$$

$$= k^2 - \frac{h^2}{12}\sum_{r=1}^{3}k_r^4 + \frac{h^4}{360}\sum_{r=1}^{3}k_r^6 + \mathcal{O}(h^6). \tag{3.33}$$

Accordingly,

$$\mu_{\lambda,h}'' + \left(\widehat{k}_h^2 - \frac{a''}{a}\right)\mu_{\lambda,h} = 0. \tag{3.34}$$

Radiation era:

$$a(\eta) \propto \eta, \qquad a'' = 0, \tag{3.35}$$

$$\mu_\kappa(\eta) = \frac{e^{-i\kappa\eta}}{\sqrt{2\kappa}}, \qquad h_\kappa = \frac{2\mu_\kappa}{aM_{\rm Pl}}. \tag{3.36}$$

Thus

$$a|h_\kappa| = \text{constant}. \tag{3.37}$$

Matter era:

$$a(\eta) \propto \eta^2, \tag{3.38}$$

$$h_\kappa(\eta) = 3\frac{j_1(\kappa\eta)}{\kappa\eta}, \tag{3.39}$$

$$h_\kappa'' + \frac{4}{\eta}h_\kappa' + \kappa^2 h_\kappa = 0. \tag{3.40}$$

The numerical orders are

$$p_T^{\text{radiation}} = 1.999618, \tag{3.41}$$
$$p_{T,\mathrm{R}}^{\text{radiation}} = 3.998798, \tag{3.42}$$
$$p_T^{\text{matter}} = 2.002687, \tag{3.43}$$
$$p_{T,\mathrm{R}}^{\text{matter}} = 3.996376. \tag{3.44}$$

The canonical Wronskian and amplitude residuals satisfy

$$\varepsilon_{\mathrm{W}} = 4.440892 \times 10^{-16}, \tag{3.45}$$
$$\varepsilon_{a|h|} = 4.440892 \times 10^{-16}. \tag{3.46}$$

3.5 Propagation distance

For geometric-optics tensor propagation, the canonical variable obeys the locally flat radial wave law

$$\mu \sim \frac{\mathrm{e}^{-\mathrm{i}k\eta}}{r}. \tag{3.47}$$

Using $h = 2\mu/(aM_{\mathrm{Pl}})$ then gives

$$h \sim \frac{1}{ar}, \tag{3.48}$$

up to the source normalization and polarization amplitude. After source-frame normalization,

$$D_L^{\mathrm{GW}} = D_L^{\mathrm{EM}} \tag{3.49}$$

in the unmodified propagation sector. Source generation, binary calibration, and modified-friction alternatives are not derived here.

3.6 Open extensions

- Add the matter-sourced vector constraint and derive the complete vorticity hierarchy.
- Include cosmological anisotropic stress and its contribution to vector and tensor propagation.
- Derive the primordial tensor normalization from the selected quantum initial state.
- Compute the leading nonlinear scalar–vector–tensor couplings.
- Incorporate source generation, binary calibration, and standard-siren observables in the tensor sector.

CHAPTER 4

Linear Matter Growth and First Structure Observables

Linear density growth supplies the first late-time structure observable. The continuum growth equation and its projector discretization are compared in Einstein–de Sitter and flat ΛCDM backgrounds.

4.1 Subhorizon dust system

For pressureless matter,

$$\dot{\delta} = -\frac{\theta}{a}, \tag{4.1}$$

$$\dot{\theta} + H\theta = -\frac{1}{a}\nabla^2\Phi, \tag{4.2}$$

$$\nabla^2\Phi = 4\pi G a^2 \bar{\rho}_m \delta. \tag{4.3}$$

Differentiate the continuity equation:

$$\ddot{\delta} = -\frac{\dot{\theta}}{a} + \frac{H\theta}{a}. \tag{4.4}$$

Using the Euler equation in the form

$$\dot{\theta} = -H\theta - \frac{1}{a}\nabla^2\Phi \tag{4.5}$$

gives

$$\ddot{\delta} = \frac{H\theta}{a} + \frac{1}{a^2}\nabla^2\Phi + \frac{H\theta}{a} \tag{4.6}$$

$$= -2H\dot{\delta} + \frac{1}{a^2}\nabla^2\Phi. \tag{4.7}$$

The Poisson equation then yields

$$\frac{1}{a^2}\nabla^2\Phi = 4\pi G\bar{\rho}_m\delta, \tag{4.8}$$

so elimination of θ and Φ gives

$$\boxed{\ddot{\delta} + 2H\dot{\delta} - 4\pi G\bar{\rho}_m\delta = 0.} \tag{4.9}$$

For the growth factor $D(a)$,

$$\dot{D} = aHD_{,a}, \tag{4.10}$$

$$\ddot{D} = a^2H^2D_{,aa} + \left(aH^2 + a^2HH_{,a}\right)D_{,a}. \tag{4.11}$$

Substitution into the time-domain growth equation gives

$$a^2H^2D_{,aa} + \left(3aH^2 + a^2HH_{,a}\right)D_{,a} - 4\pi G\bar{\rho}_m D = 0. \tag{4.12}$$

Using

$$4\pi G\bar{\rho}_m = \frac{3}{2}\Omega_{m0}H_0^2 a^{-3} \tag{4.13}$$

and dividing by a^2H^2 yields

$$\boxed{D_{,aa} + \left(\frac{3}{a} + \frac{H_{,a}}{H}\right)D_{,a} - \frac{3\Omega_{m0}H_0^2}{2a^5H(a)^2}D = 0.} \tag{4.14}$$

With $N = \log a$,

$$D_{,a} = \frac{1}{a}D_{,N}, \tag{4.15}$$

$$D_{,aa} = \frac{1}{a^2}\left(D_{,NN} - D_{,N}\right), \tag{4.16}$$

$$\frac{aH_{,a}}{H} = \frac{\mathrm{d}\log H}{\mathrm{d}N}, \tag{4.17}$$

$$\Omega_m(a) = \frac{\Omega_{m0}H_0^2a^{-3}}{H(a)^2}. \tag{4.18}$$

Multiplying the a-equation by a^2 and inserting these identities gives

$$\boxed{D_{,NN} + \left[2 + \frac{\mathrm{d}\log H}{\mathrm{d}N}\right]D_{,N} - \frac{3}{2}\Omega_m(a)D = 0.} \tag{4.19}$$

4.2 Einstein–de Sitter benchmark

For

$$\Omega_m(a) = 1, \qquad \frac{\mathrm{d}\log H}{\mathrm{d}N} = -\frac{3}{2}, \tag{4.20}$$

a power law $D \propto a^r = \mathrm{e}^{rN}$ gives the characteristic equation

$$r^2 + \frac{1}{2}r - \frac{3}{2} = 0, \tag{4.21}$$

$$(r-1)\left(r + \frac{3}{2}\right) = 0. \tag{4.22}$$

Hence

$$\boxed{D_+(a) = a, \qquad D_-(a) = a^{-3/2}.} \tag{4.23}$$

The midpoint solver converges with

$$p_{\mathrm{EdS}} = 1.968111. \tag{4.24}$$

4.3 Flat ΛCDM growth

Let

$$E(a)^2 = \Omega_{m0}a^{-3} + 1 - \Omega_{m0}, \tag{4.25}$$

$$\Omega_m(a) = \frac{\Omega_{m0}a^{-3}}{E(a)^2}. \tag{4.26}$$

Define

$$I(a) := \int_0^a \frac{\mathrm{d}a'}{a'^3E(a')^3}, \qquad I_{,a} = \frac{1}{a^3E(a)^3}. \tag{4.27}$$

Reduction of order about the decaying solution proportional to $E(a)$ gives a second solution proportional to $E(a)I(a)$. Normalizing it to the matter-era behavior $D_+ \sim a$ fixes the coefficient, so

$$\boxed{D_+(a) = \frac{5\Omega_{m0}}{2}E(a)\int_0^a \frac{\mathrm{d}a'}{a'^3E(a')^3}.} \tag{4.28}$$

For $\Omega_{m0} = 0.3$, the growth ODE and integral solution agree with trajectory residual

$$\varepsilon_{\mathrm{Heath}} = 1.035117 \times 10^{-12}. \tag{4.29}$$

The midpoint convergence order is

$$p_{\Lambda\mathrm{CDM}} = 1.970169. \tag{4.30}$$

4.4 Consistent projector Laplacian

Let

$$\Delta_h \mathrm{e}^{\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}} = -\widehat{k}_h^2 \mathrm{e}^{\mathrm{i}\boldsymbol{k}\cdot\boldsymbol{x}}. \tag{4.31}$$

The lattice Poisson equation is

$$-\widehat{k}_h^2\Phi_h = 4\pi G a^2 \bar{\rho}_m \delta_h. \tag{4.32}$$

The lattice Euler force is

$$\dot{\theta}_h + H\theta_h = \frac{\widehat{k}_h^2}{a}\Phi_h. \tag{4.33}$$

Therefore

$$\boxed{\frac{\widehat{k}_h^2}{a}\Phi_h = -4\pi G a \bar{\rho}_m \delta_h.} \tag{4.34}$$

Differentiating $\dot{\delta}_h = -\theta_h/a$ gives

$$\ddot{\delta}_h = -\frac{\dot{\theta}_h}{a} + \frac{H\theta_h}{a} \tag{4.35}$$

$$= -2H\dot{\delta}_h - \frac{\widehat{k}_h^2}{a^2}\Phi_h. \tag{4.36}$$

The discrete Poisson relation now removes the same symbol that appears in the Euler force:

$$-\frac{\widehat{k}_h^2}{a^2}\Phi_h = 4\pi G\bar{\rho}_m\delta_h. \tag{4.37}$$

Hence the common projector Laplacian cancels exactly:

$$\boxed{\ddot{\delta}_h + 2H\dot{\delta}_h - 4\pi G\bar{\rho}_m\delta_h = 0.} \tag{4.38}$$

Thus a consistent Euler–Poisson pair introduces no artificial scale dependence in pressureless linear growth.

4.5 Growth observables

Define

$$f(a) := \frac{\mathrm{d}\log D}{\mathrm{d}\log a} = \frac{(1/D)\,\mathrm{d}D}{\mathrm{d}N} = \frac{D_{,N}}{D}, \tag{4.39}$$

$$D_N(a) := \frac{D(a)}{D(1)}, \tag{4.40}$$

$$\sigma_8(a) = D_N(a)\sigma_{8,0}, \tag{4.41}$$

$$f\sigma_8(a) := f(a)\sigma_8(a) = f(a)D_N(a)\sigma_{8,0}. \tag{4.42}$$

For

$$\Omega_{m0} = 0.3, \qquad \sigma_{8,0} = 0.8, \tag{4.43}$$

the controlled benchmark gives

$$f(1) = 0.512796, \tag{4.44}$$

$$f\sigma_8(1) = 0.410237. \tag{4.45}$$

These values are conditional on the stated background parameters and normalization, not an observational fit.

4.6 Open extensions

- Add radiation, baryon, and neutrino transfer functions to the projector growth system.
- Extend the evolution through the superhorizon regime and include free-streaming species.
- Normalize the matter power spectrum consistently from the primordial sector.
- Map the predicted growth history to survey observables and a reproducible likelihood.
- Develop the nonlinear fluid or N-body limit while preserving the projector constraints.

CHAPTER 5

Controlled Primordial Power-Law Sector

A constant slow-roll parameter gives a solvable inflationary benchmark. The scalar mode equation and power spectra are used here as a controlled test of the projector discretization.

5.1 Constant-ϵ background

Assume

$$0 < \epsilon < 1, \qquad q := \frac{1}{1-\epsilon}, \qquad \eta < 0, \tag{5.1}$$

and

$$\boxed{a(\eta) = A(-\eta)^{-q}.} \tag{5.2}$$

Differentiation gives

$$a' = Aq(-\eta)^{-q-1}, \tag{5.3}$$

$$\mathcal{H} = \frac{a'}{a} = -\frac{q}{\eta}, \tag{5.4}$$

$$\mathcal{H}' = \frac{q}{\eta^2}. \tag{5.5}$$

Hence

$$1 - \frac{\mathcal{H}'}{\mathcal{H}^2} = 1 - \frac{1}{q} = \epsilon. \tag{5.6}$$

Then

$$\mathcal{H} := \frac{a'}{a} = -\frac{q}{\eta}, \tag{5.7}$$

$$\epsilon_H := 1 - \frac{\mathcal{H}'}{\mathcal{H}^2} = \epsilon. \tag{5.8}$$

The comoving Hubble radius obeys

$$(aH)^{-1} = \mathcal{H}^{-1} = -\frac{\eta}{q}. \tag{5.9}$$

For $N := \log(a/a_i)$,

$$\frac{(aH)^{-1}}{(a_i H_i)^{-1}} = \mathrm{e}^{-(1-\epsilon)N}, \tag{5.10}$$

$$\frac{|\Omega_K|}{|\Omega_{K,i}|} = \mathrm{e}^{-2(1-\epsilon)N}. \tag{5.11}$$

For constant $\epsilon = -\,\mathrm{d}\log H/\,\mathrm{d}N$,

$$H/H_i = \mathrm{e}^{-\epsilon N}, \qquad a/a_i = \mathrm{e}^{N}. \tag{5.12}$$

Therefore $aH/(a_iH_i) = \mathrm{e}^{(1-\epsilon)N}$. Since $|\Omega_K| \propto (aH)^{-2}$, the curvature ratio is the square of the comoving-Hubble-radius ratio. Hence

$$N \geq \frac{\log(S_K^{-1})}{2(1-\epsilon)} \Longrightarrow \frac{|\Omega_K|}{|\Omega_{K,i}|} \leq S_K. \tag{5.13}$$

5.2 Canonical scalar mode

For a canonical single field with constant ϵ,

$$z := a\sqrt{2\epsilon}\, M_{\mathrm{Pl}}, \qquad \nu := q + \frac{1}{2} = \frac{3-\epsilon}{2(1-\epsilon)}. \tag{5.14}$$

Then

$$\boxed{\frac{z''}{z} = \frac{a''}{a} = \frac{\nu^2 - \frac{1}{4}}{\eta^2}.} \tag{5.15}$$

Because ϵ is constant, z differs from a only by a constant factor. Also,

$$\frac{a''}{a} = \frac{q(q+1)}{\eta^2}, \tag{5.16}$$

$$\nu^2 - \frac{1}{4} = \left(q + \frac{1}{2}\right)^2 - \frac{1}{4} \tag{5.17}$$

$$= q(q+1). \tag{5.18}$$

The quadratic mode action is

$$S_v^{(2)} = \frac{1}{2}\int \mathrm{d}\eta\, \mathrm{d}^3k \left[|v_k'|^2 - \left(k^2 - \frac{z''}{z}\right)|v_k|^2\right]. \tag{5.19}$$

Varying v_k^* and integrating the kinetic term by parts gives

$$\delta S_v^{(2)} = -\frac{1}{2}\int \delta v_k^* \left[v_k'' + \left(k^2 - \frac{z''}{z}\right) v_k\right] \mathrm{d}\eta\, \mathrm{d}^3k + \text{c.c.}, \tag{5.20}$$

so stationarity yields the mode equation. Thus

$$\boxed{v_k'' + \left[k^2 - \frac{\nu^2 - \frac{1}{4}}{\eta^2}\right] v_k = 0.} \tag{5.21}$$

The Bunch–Davies fixed-background mode is

$$\boxed{v_k(\eta) = \frac{\sqrt{-\pi\eta}}{2}\mathrm{e}^{\mathrm{i}(2\nu+1)\pi/4} H_\nu^{(1)}(-k\eta).} \tag{5.22}$$

Its normalization is

$$v_k v_k^{*\prime} - v_k^* v_k' = \mathrm{i}. \tag{5.23}$$

The vacuum choice and fixed-background canonical quantization are model inputs in this section; quantized projector geometry is not derived here.

5.3 Exact power-law spectra

For $a = A(-\eta)^{-q}$, the superhorizon scalar spectrum is

$$\boxed{\mathcal{P}_{\mathcal{R}}(k) = \frac{2^{2\nu-4}\Gamma(\nu)^2}{\pi^3 A^2 \epsilon M_{\mathrm{Pl}}^2} k^{3-2\nu}.} \tag{5.24}$$

For $-k\eta \ll 1$,

$$H_\nu^{(1)}(x) \sim -\frac{\mathrm{i}\Gamma(\nu)}{\pi}\left(\frac{2}{x}\right)^\nu. \tag{5.25}$$

Thus

$$|v_k|^2 \sim \frac{2^{2\nu-2}\Gamma(\nu)^2}{\pi}(-\eta)^{1-2\nu}k^{-2\nu}. \tag{5.26}$$

Since $\mathcal{R}_k = v_k/z$, $z = A(-\eta)^{-q}\sqrt{2\epsilon}M_{\mathrm{Pl}}$, and $\nu = q + 1/2$, all powers of η cancel. Finally,

$$\mathcal{P}_{\mathcal{R}} = \frac{k^3}{2\pi^2}|\mathcal{R}_k|^2, \tag{5.27}$$

which gives the displayed coefficient and power of k. The two-polarization tensor spectrum is

$$\boxed{\mathcal{P}_T(k) = \frac{2^{2\nu}\Gamma(\nu)^2}{\pi^3 A^2 M_{\mathrm{Pl}}^2} k^{3-2\nu}.} \tag{5.28}$$

Therefore

$$n_s - 1 = 3 - 2\nu = -\frac{2\epsilon}{1-\epsilon}, \tag{5.29}$$

$$n_t = -\frac{2\epsilon}{1-\epsilon}, \tag{5.30}$$

$$\alpha_s := \frac{\mathrm{d}n_s}{\mathrm{d}\log k} = 0, \tag{5.31}$$

$$\alpha_t := \frac{\mathrm{d}n_t}{\mathrm{d}\log k} = 0, \tag{5.32}$$

$$\boxed{r := \frac{\mathcal{P}_T}{\mathcal{P}_{\mathcal{R}}} = 16\epsilon}, \tag{5.33}$$

$$\boxed{r = -8(1-\epsilon)n_t}. \tag{5.34}$$

The spectral index is the logarithmic derivative of the power law:

$$\frac{\mathrm{d}\log \mathcal{P}_{\mathcal{R}}}{\mathrm{d}\log k} = 3 - 2\nu. \tag{5.35}$$

Using $\nu = (3-\epsilon)/[2(1-\epsilon)]$ gives $3 - 2\nu = -2\epsilon/(1-\epsilon)$. The scalar and tensor spectra have the same k-power, so their runnings vanish. Dividing their normalization constants gives

$$r = \frac{2^{2\nu}}{2^{2\nu-4}/\epsilon} = 16\epsilon, \tag{5.36}$$

and substitution of n_t gives the exact consistency relation. The first-order slow-roll relation $r = -8n_t + O(\epsilon^2)$ follows.

5.4 Controlled benchmark

For

$$\epsilon = 0.01, \qquad N = 60, \tag{5.37}$$

The power-law calculation gives

$$q = 1.010101, \tag{5.38}$$
$$\nu = 1.510101, \tag{5.39}$$
$$n_s = 0.979798, \tag{5.40}$$
$$n_t = -0.020202, \tag{5.41}$$
$$r = 0.16, \tag{5.42}$$
$$\frac{|\Omega_K|}{|\Omega_{K,i}|} = 2.545749 \times 10^{-52}, \tag{5.43}$$
$$\frac{(aH)^{-1}}{(a_i H_i)^{-1}} = 1.595540 \times 10^{-26}. \tag{5.44}$$

The numerical checks are

$$\varepsilon_{\epsilon_H} = 0, \tag{5.45}$$
$$\varepsilon_{z''/z} = 0, \tag{5.46}$$
$$\varepsilon_{n_s} = 1.179612 \times 10^{-16}, \tag{5.47}$$
$$\varepsilon_{n_t} = 2.255141 \times 10^{-16}, \tag{5.48}$$
$$\varepsilon_r = 2.775558 \times 10^{-17}, \tag{5.49}$$
$$\varepsilon_{\mathrm{W}} = 3.907985 \times 10^{-14}, \tag{5.50}$$
$$p_{v,\mathrm{mid}} = 1.964609. \tag{5.51}$$

5.5 Open extensions

- Derive the conditions that select the inflationary attractor from general initial data.
- Construct a graceful-exit mechanism and a controlled reheating sector.
- Determine the vacuum from projector refinement instead of fixing it by hand.
- Quantize the projector geometry together with the scalar perturbations.
- Extend the calculation to multifield dynamics and the leading non-Gaussian observables.

CHAPTER 6

Recombination, Decoupling, and Acoustic Scales

Recombination and acoustic scales are introduced at the level needed to connect the background and perturbation sectors to observable angular scales. Atomic rates and tight-coupling formulas are treated as standard model inputs.

6.1 Conditional atomic benchmark

The benchmark inputs are

$$H_0 = 70\ \mathrm{km\,s^{-1}\,Mpc^{-1}}, \tag{6.1}$$

$$\Omega_b = 0.05, \qquad \Omega_m = 0.30, \tag{6.2}$$

$$\Omega_r = 9 \times 10^{-5}, \qquad T_0 = 2.7255\ \mathrm{K}, \tag{6.3}$$

$$Y_p = 0.24. \tag{6.4}$$

These values define a reproducible controlled model; they are not an observational fit.

Let

$$T(z) = T_0(1+z), \tag{6.5}$$

$$n_H(z) = n_{H0}(1+z)^3, \tag{6.6}$$

$$S(z) = \frac{1}{n_H(z)} \left(\frac{m_e k_B T(z)}{2\pi\hbar^2} \right)^{3/2} \exp\left[-\frac{\chi_H}{k_B T(z)} \right]. \tag{6.7}$$

The Saha equilibrium equation is

$$\frac{x_e^2}{1-x_e} = S. \tag{6.8}$$

Multiplication by $1 - x_e$ gives the quadratic equation

$$x_e^2 + S x_e - S = 0. \tag{6.9}$$

Hence

$$x_e = \frac{-S \pm \sqrt{S^2 + 4S}}{2}. \tag{6.10}$$

For $S > 0$, the minus sign is negative, so the physical root is the plus branch. Rationalizing that branch gives

$$x_e(S) = \frac{-S + \sqrt{S^2 + 4S}}{2} \tag{6.11}$$

$$= \frac{2S}{S + \sqrt{S^2 + 4S}} \tag{6.12}$$

$$= \frac{2}{1 + \sqrt{1 + 4/S}}. \tag{6.13}$$

Thus

$$\boxed{x_e(S) = \frac{-S + \sqrt{S^2 + 4S}}{2} = \frac{2}{1 + \sqrt{1 + 4/S}}.} \tag{6.14}$$

6.2 Optical depth and visibility

Ignoring reionization,

$$n_e(z) = x_e(z) n_H(z), \tag{6.15}$$

$$\frac{\mathrm{d}\tau}{\mathrm{d}z} = \frac{c\sigma_T n_e(z)}{(1+z)H(z)}, \tag{6.16}$$

$$\tau(z) = \int_0^z \frac{c\sigma_T n_e(z')}{(1+z')H(z')}\,\mathrm{d}z'. \tag{6.17}$$

The survival probability from redshift z to the present is $\mathrm{e}^{-\tau(z)}$. Its loss per unit redshift is

$$-\frac{\mathrm{d}}{\mathrm{d}z}\mathrm{e}^{-\tau(z)} = -\left[-\mathrm{e}^{-\tau(z)}\frac{\mathrm{d}\tau}{\mathrm{d}z}\right] \tag{6.18}$$

$$= \mathrm{e}^{-\tau(z)}\frac{\mathrm{d}\tau}{\mathrm{d}z}. \tag{6.19}$$

The redshift-space visibility is therefore

$$\boxed{g_z(z) = \mathrm{e}^{-\tau(z)}\frac{\mathrm{d}\tau}{\mathrm{d}z}.} \tag{6.20}$$

Integrating the total derivative gives

$$\int_0^{z_{\max}} g_z(z)\,\mathrm{d}z = -\int_0^{z_{\max}} \frac{\mathrm{d}}{\mathrm{d}z}\mathrm{e}^{-\tau(z)}\,\mathrm{d}z \tag{6.21}$$

$$= \mathrm{e}^{-\tau(0)} - \mathrm{e}^{-\tau(z_{\max})} \tag{6.22}$$

$$= 1 - \mathrm{e}^{-\tau(z_{\max})}. \tag{6.23}$$

The controlled Saha benchmark gives

$$z_* = 1282.15, \tag{6.24}$$

$$z_{\tau=1} = 1289.423097, \tag{6.25}$$

$$\Delta z_{\mathrm{FWHM}} = 139.2. \tag{6.26}$$

The visibility-normalization residual is

$$\varepsilon_g = 4.807365 \times 10^{-8}. \tag{6.27}$$

Nonequilibrium multilevel recombination and reionization are not derived here.

6.3 Baryon loading and sound horizon

Since nonrelativistic baryons and photons scale as

$$\rho_b(z) = \rho_{b0}(1+z)^3, \qquad \rho_\gamma(z) = \rho_{\gamma 0}(1+z)^4, \tag{6.28}$$

the loading factor is

$$R(z) := \frac{3\rho_b}{4\rho_\gamma} \tag{6.29}$$

$$= \frac{3\rho_{b0}(1+z)^3}{4\rho_{\gamma 0}(1+z)^4} \tag{6.30}$$

$$= \frac{3\Omega_b}{4\Omega_\gamma(1+z)}, \tag{6.31}$$

$$c_s(z) := \frac{c}{\sqrt{3(1+R(z))}}. \tag{6.32}$$

The comoving distance travelled by an acoustic signal before decoupling is

$$r_s(t_*) = \int_0^{t_*} \frac{c_s(t)}{a(t)}\,\mathrm{d}t. \tag{6.33}$$

Using

$$\frac{\mathrm{d}z}{\mathrm{d}t} = -(1+z)H(z), \qquad a = (1+z)^{-1}, \tag{6.34}$$

one has

$$\frac{\mathrm{d}t}{a} = -\frac{\mathrm{d}z}{H(z)}. \tag{6.35}$$

Reversing the limits therefore gives

$$\boxed{r_s(z_*) = \int_{z_*}^{\infty} \frac{c_s(z)}{H(z)}\,\mathrm{d}z.} \tag{6.36}$$

The line-of-sight comoving distance is

$$\chi_* := \int_0^{z_*} \frac{c}{H(z)}\,\mathrm{d}z. \tag{6.37}$$

The geometric acoustic multipole is

$$\boxed{\ell_A := \pi \frac{\chi_*}{r_s(z_*)}.} \tag{6.38}$$

For the benchmark,

$$r_s(z_*) = 125.784218\ \mathrm{Mpc}, \tag{6.39}$$

$$\chi_* = 13644.312614\ \mathrm{Mpc}, \tag{6.40}$$

$$\ell_A = 340.781006. \tag{6.41}$$

This ℓ_A is a geometric acoustic ruler, not a claim that the first CMB peak equals ℓ_A.

6.4 Tight-coupling oscillator

For constant loading and absent gravitational driving,

$$X'' + k^2 c_s^2 X = 0. \tag{6.42}$$

Writing $\omega_s := kc_s$, the characteristic roots are $r = \pm i\omega_s$, so

$$X(\eta) = A\cos(\omega_s \Delta\eta) + B\sin(\omega_s \Delta\eta). \tag{6.43}$$

At the initial time, $A = X_i$ and $X_i' = \omega_s B$. Therefore

$$X(\eta) = X_i \cos(kc_s \Delta\eta) + \frac{X_i'}{kc_s}\sin(kc_s \Delta\eta). \tag{6.44}$$

For varying $R(a)$, the acoustic model integrates

$$X'' + k^2 c_s(a)^2 X = 0 \tag{6.45}$$

through the benchmark background. The midpoint refinement order is

$$p_{\text{acoustic}} = 2.021414. \tag{6.46}$$

The final-state error at the finest step is

$$\varepsilon_{\text{acoustic}} = 3.978519 \times 10^{-5}. \tag{6.47}$$

6.5 Diffusion approximation

Let the positive conformal opacity be

$$\kappa' := a n_e \sigma_T. \tag{6.48}$$

The controlled tight-coupling diffusion estimate is

$$\boxed{\frac{1}{k_D^2(\eta_*)} = \int_0^{\eta_*} \frac{\mathrm{d}\eta}{6\kappa'} \frac{R^2 + \frac{16}{15}(1+R)}{(1+R)^2}.} \tag{6.49}$$

At leading diffusion order, each Fourier amplitude obeys a local attenuation equation of the form

$$\frac{\mathrm{d}X_k}{\mathrm{d}\eta} = -k^2 D_\gamma(\eta) X_k, \tag{6.50}$$

where the tight-coupling coefficient integrates to

$$\int_0^{\eta_*} D_\gamma(\eta)\,\mathrm{d}\eta = \frac{1}{k_D^2(\eta_*)}. \tag{6.51}$$

Separating variables gives

$$\frac{\mathrm{d}X_k}{X_k} = -k^2 D_\gamma \,\mathrm{d}\eta, \tag{6.52}$$

$$\log \frac{X_k(\eta_*)}{X_k(0)} = -\frac{k^2}{k_D^2}, \tag{6.53}$$

so the associated damping factor is

$$\mathcal{D}(k) = \exp\left[-\frac{k^2}{k_D^2}\right]. \tag{6.54}$$

The benchmark gives

$$r_D := k_D^{-1} = 4.554719\ \mathrm{Mpc}, \tag{6.55}$$

$$k_D = 0.219553\ \mathrm{Mpc}^{-1}. \tag{6.56}$$

The diffusion formula is a tight-coupling approximation and does not replace a full collision hierarchy.

6.6 Open extensions

- Replace Saha equilibrium with a multilevel nonequilibrium hydrogen-recombination calculation.
- Add helium recombination and a controlled reionization history.
- Include neutrino decoupling and its effect on the acoustic background.
- Derive the full photon collision hierarchy beyond the tight-coupling approximation.
- Improve diffusion damping and calibrate the drag epoch and sound horizon at precision accuracy.

CHAPTER 7

Cosmological Finite-Scale Photon Propagation

The finite-spacing photon correction is integrated along an expanding line of sight. The result distinguishes physical and comoving microscopic scales and yields a conditional arrival-time formula.

7.1 Local velocity law

For the selected cubic microscopic walk,

$$\cos(\omega\ell) = \prod_{i=1}^{3} \cos(kn_i\ell), \tag{7.1}$$

$$\mathcal{A}(\boldsymbol{n}) := \sum_{i<j} n_i^2 n_j^2, \tag{7.2}$$

$$v_g(k, \boldsymbol{n}, \ell) = 1 - \frac{1}{2}\mathcal{A}(\boldsymbol{n})\ell^2 k^2 + \mathcal{O}(\ell^4). \tag{7.3}$$

The exact radial group velocity is

$$v_g = \frac{\left[\prod_i \cos(kn_i\ell)\right] \sum_i n_i \tan(kn_i\ell)}{\sin(\omega\ell)}. \tag{7.4}$$

7.2 Physical versus comoving microscopic scale

Let

$$E(z) = E_0(1+z), \tag{7.5}$$

$$\ell_{\mathrm{phys}}(z) = \ell_{*,0}(1+z)^{-\sigma}. \tag{7.6}$$

Then

$$\ell_{\mathrm{phys}}(z)E(z) = \ell_{*,0}(1+z)^{-\sigma}E_0(1+z) \tag{7.7}$$

$$= \ell_{*,0}E_0(1+z)^{1-\sigma}, \tag{7.8}$$

and therefore

$$\left[\frac{\ell_{\mathrm{phys}}(z)E(z)}{\hbar c}\right]^2 = \left[\frac{\ell_{*,0}E_0}{\hbar c}\right]^2 (1+z)^{2-2\sigma}. \tag{7.9}$$

The cases

$$\sigma = 0 \qquad \Longleftrightarrow \text{ fixed proper microscopic length,} \tag{7.10}$$

$$\sigma = 1 \qquad \Longleftrightarrow \text{ fixed comoving microscopic length} \tag{7.11}$$

are physically distinct.

7.3 Exact arrival condition

Let the lower-energy photon arrive at $z = 0$. Its source comoving distance is

$$\chi_s = \int_0^{z_s} \frac{v_\ell(z)}{H(z)}\,\mathrm{d}z. \tag{7.12}$$

If the higher-energy photon arrives later at $z = z_{a,h} < 0$, then

$$\chi_s = \int_{z_{a,h}}^{z_s} \frac{v_h(z)}{H(z)}\,\mathrm{d}z. \tag{7.13}$$

Splitting the second integral at the present epoch and equating the two expressions for χ_s gives

$$\int_0^{z_s} \frac{v_\ell}{H}\,\mathrm{d}z = \int_{z_{a,h}}^{0} \frac{v_h}{H}\,\mathrm{d}z + \int_0^{z_s} \frac{v_h}{H}\,\mathrm{d}z, \tag{7.14}$$

$$\int_{z_{a,h}}^{0} \frac{v_h}{H}\,\mathrm{d}z = \int_0^{z_s} \frac{v_\ell - v_h}{H}\,\mathrm{d}z. \tag{7.15}$$

The exact background arrival-time difference is

$$\Delta t_{h\ell}^{\mathrm{direct}} = \int_{z_{a,h}}^{0} \frac{\mathrm{d}z}{(1+z)H(z)}. \tag{7.16}$$

7.4 Leading cosmological delay law

Define

$$I_\sigma(z_s) := \int_0^{z_s} \frac{(1+z)^{2-2\sigma}}{H(z)}\,\mathrm{d}z. \tag{7.17}$$

Writing $k = E/(\hbar c)$, the local velocity deficit is

$$1 - v_g = \frac{1}{2}\mathcal{A}(\boldsymbol{n})\ell_{\mathrm{phys}}(z)^2\frac{E(z)^2}{\hbar^2 c^2} + \mathcal{O}(\ell_{*,0}^4) \tag{7.18}$$

$$= \frac{1}{2}\mathcal{A}(\boldsymbol{n})\frac{\ell_{*,0}^2 E_0^2}{\hbar^2 c^2}(1+z)^{2-2\sigma} + \mathcal{O}(\ell_{*,0}^4). \tag{7.19}$$

Hence

$$v_\ell - v_h = (1 - v_h) - (1 - v_\ell) \tag{7.20}$$

$$= \frac{\mathcal{A}(\boldsymbol{n})\ell_{*,0}^2}{2\hbar^2 c^2}(E_{h,0}^2 - E_{\ell,0}^2)(1+z)^{2-2\sigma} + \mathcal{O}(\ell_{*,0}^4). \tag{7.21}$$

The arrival redshift satisfies $z_{a,h} = \mathcal{O}(\ell_{*,0}^2)$. On the short interval from $z_{a,h}$ to zero, replacing v_h and $(1+z)^{-1}$ by one changes the result only at order $\ell_{*,0}^4$. The distance-balance identity therefore becomes

$$\Delta t_{h\ell} = \int_0^{z_s} \frac{v_\ell - v_h}{H(z)}\,\mathrm{d}z + \mathcal{O}(\ell_{*,0}^4). \tag{7.22}$$

Thus, to leading order in $\ell_{*,0}^2$,

$$\boxed{\Delta t_{h\ell} = \frac{\mathcal{A}(\boldsymbol{n})\ell_{*,0}^2(E_{h,0}^2 - E_{\ell,0}^2)}{2\hbar^2c^2} I_\sigma(z_s) + \mathcal{O}(\ell_{*,0}^4).} \tag{7.23}$$

Special cases:

$$\sigma = 0: \quad \Delta t_{h\ell}^{\text{proper}} = \frac{\mathcal{A}\ell_{*,0}^2\Delta E_0^2}{2\hbar^2c^2}\int_0^{z_s}\frac{(1+z)^2}{H(z)}\,\mathrm{d}z + \mathcal{O}(\ell_{*,0}^4), \tag{7.24}$$

$$\sigma = 1: \quad \Delta t_{h\ell}^{\text{comoving}} = \frac{\mathcal{A}\ell_{*,0}^2\Delta E_0^2}{2\hbar^2c^2}\int_0^{z_s}\frac{\mathrm{d}z}{H(z)} + \mathcal{O}(\ell_{*,0}^4). \tag{7.25}$$

At low redshift,

$$H(z) = H_0 + \mathcal{O}(z), \tag{7.26}$$

$$(1+z)^{2-2\sigma} = 1 + \mathcal{O}(z), \tag{7.27}$$

$$I_\sigma(z_s) = \frac{z_s}{H_0} + \mathcal{O}(z_s^2). \tag{7.28}$$

Since $L = cz_s/H_0 + \mathcal{O}(z_s^2)$,

$$I_\sigma(z_s) = \frac{L}{c} + \mathcal{O}(z_s^2), \tag{7.29}$$

which recovers the static-distance law.

7.5 Null bound and directional cancellation

A null result $|\Delta t_{h\ell}| \le \Delta t_{\max}$ implies

$$\frac{\mathcal{A}(\boldsymbol{n})\ell_{*,0}^2|E_{h,0}^2 - E_{\ell,0}^2|}{2\hbar^2c^2} I_\sigma(z_s) \le \Delta t_{\max}. \tag{7.30}$$

Solving this inequality for the positive length $\ell_{*,0}$ gives

$$\boxed{\ell_{*,0} \le \left[\frac{2\hbar^2c^2\Delta t_{\max}}{\mathcal{A}(\boldsymbol{n})|E_{h,0}^2 - E_{\ell,0}^2| I_\sigma(z_s)}\right]^{1/2}.} \tag{7.31}$$

For a cubic-axis ray,

$$\mathcal{A}(1,0,0) = 0. \tag{7.32}$$

The exact dispersion also reduces to

$$\cos(\omega\ell) = \cos(k\ell)\cos(0)\cos(0) \tag{7.33}$$

$$= \cos(k\ell). \tag{7.34}$$

On the principal branch, $\omega = k$, and hence

$$v_g = \frac{\partial \omega}{\partial k} = 1. \tag{7.35}$$

Thus the leading and exact directional delays vanish on the axis.

7.6 Numerical verification

For

$$H(z) = H_0\sqrt{0.3(1+z)^3 + 0.7}, \qquad z_s = 2, \tag{7.36}$$

direct arrival integration gives

$$p^{\sigma=0}_{\text{remainder}} = 4.002683, \tag{7.37}$$
$$p^{\sigma=1}_{\text{remainder}} = 4.001430, \tag{7.38}$$
$$\varepsilon_{\text{axis}} = 0, \tag{7.39}$$
$$\ell_{\text{true}} = 0.030000, \tag{7.40}$$
$$\ell^{\sigma=0}_{\text{fit}} = 0.030008, \tag{7.41}$$
$$\ell^{\sigma=1}_{\text{fit}} = 0.030003. \tag{7.42}$$

Hence

$$\Delta t^{\text{direct}}_{h\ell} - \Delta t^{\text{leading}}_{h\ell} = \mathcal{O}(\ell^4_{*,0}). \tag{7.43}$$

7.7 Open extensions

- Derive the finite-scale correction for a broader class of microscopic walks and determine which terms are universal.
- Remove the preferred cubic frame, or derive and quantify the residual frame dependence after refinement.
- Model intrinsic source lags and detector systematics in the same arrival-time analysis.
- Use a declared propagation data set to place a reproducible bound on the microscopic length scale.
- Derive a second independent observable that tests the same finite-scale parameter.

CHAPTER 8

Visible Averages, Hidden Detail Modes, and Model Genealogy

From geometry to a matter model. The second part separates exact fiber decomposition from proposed physical mechanism. It records the additive no-go result, defines the scale-covariant action, and derives production, coldness, locking, and release.

8.1 Scope

The cosmological construction begins with a field on spacetime and an internal refinement fiber. The fiber average is the visible sector, while orthogonal Haar-detail modes form a hidden tower. The separation is exact at the level of the fiber inner product. Its physical interpretation remains a model proposition until the refinement fiber and its couplings are derived from a more primitive microscopic theory.

8.2 Average and detail decomposition

Fix a prime p and let D_0 denote the constant Haar sector on $\mathbb{Z}_p$. Let D_h, $h \geq 1$, denote the orthogonal detail projectors, with

$$d_h = \dim \operatorname{Ran} \mathsf{D}_h = (p-1)p^{h-1}. \tag{8.1}$$

For refinement order $\sigma > 0$, define

$$\mathcal{L}_{\sigma,p} = \Lambda^2 \sum_{h\geq 1} p^{2\sigma h} \mathsf{D}_h. \tag{8.2}$$

For a real field $\Phi(x,u)$ on $M^{3,1} \times \mathbb{Z}_p$, write

$$\Phi(x,u) = \phi_0(x) + \sum_{h\geq 1} \sum_{a=1}^{d_h} \chi_{h,a}(x)\psi_{h,a}(u). \tag{8.3}$$

A quadratic fiber action decomposes into one average field and a countable tower, with representative level law

$$m_h^2 = m_0^2 + \Lambda^2 p^{2\sigma h}. \tag{8.4}$$

8.3 Hidden by averaging, not by absence

Every nonconstant Haar function has zero fiber mean, so

$$\int_{\mathbb{Z}_p} \Phi(x,u)\ \mathrm{d}\mu_p(u) = \phi_0(x). \tag{8.5}$$

If visible operators couple only to this average, detail modes have no linear visible coupling while their stress tensor remains gravitationally active:

$$\boxed{\text{zero visible average} \quad + \quad \text{nonzero gravitational stress energy.}} \tag{8.6}$$

8.4 Rejected additive branch

The first mechanism asked one light produced shell to provide the abundance, remain cold, and stabilize the average mode. The scale-covariant chapter proves these requirements incompatible on the declared one-shell, one-crossing, energy-matched branch. The rejected model is retained because it provides a genuine counterexample to an attractive but inconsistent choice.

8.5 Current branch

The current model changes the whole refinement scale through one coupling,

$$\mathcal{L}_{\sigma,p} \longmapsto \left(1 + \gamma \frac{\phi^2}{v^2}\right) \mathcal{L}_{\sigma,p}. \tag{8.7}$$

Heavier levels then acquire broader nonadiabatic scales and become colder after redshift, while their weighted energy stabilizes the average mode:

$$\boxed{\begin{aligned} \text{detail tower} &\longrightarrow \text{clustered dark matter,} \\ \text{locked average mode} &\longrightarrow \text{vacuum-like dark energy.} \end{aligned}} \tag{8.8}$$

The coupling is a proposed extension, not a theorem forced uniquely by projector algebra.

8.6 Epistemic status

The Haar decomposition and multiplicities are mathematical. Production, locking, and release follow conditionally from the declared action and state-selection approximation. Identification with the observed dark sector requires radiative stability, parameter inference, and empirical comparison.

CHAPTER 9

Scale-Covariant Refinement Dark Sector

Production hypothesis

The rapid-crossing formulas in this chapter describe a specified-background approximation. The covariant fiber analysis shows that the benchmark spectrum overdraws its available energy. Its normalized shape is therefore an input to downstream transport, not an established solution of the coupled production problem.

9.1 Scope and objective

The preceding unified dark-sector model used an additive average-detail Hessian,

$$m_h^2(\phi) = m_h^2 + 3\lambda\phi^2, \tag{9.1}$$

so that a light produced shell could both cluster and stabilize the average mode. A hard abundance and coldness audit exposes a restricted but serious obstruction in that branch. On an energy-matched rapid crossing, one dominant additive shell cannot simultaneously reproduce the present dark abundance, remain nonrelativistic, and lock the average mode.

This chapter first proves that obstruction under its stated hypotheses. It then replaces the additive Hessian by one scale-covariant deformation of the whole refinement operator,

$$\mathcal{L}_{\sigma,p} \longmapsto \left(1 + \gamma\frac{\phi^2}{v^2}\right)\mathcal{L}_{\sigma,p}. \tag{9.2}$$

The replacement introduces one dimensionless coupling γ, not an independently chosen coupling at every level. It has two linked effects:

1. heavier detail levels are created with proportionally broader nonadiabatic scales and become colder after redshift;
2. the same weighted detail energy, rather than the inverse-mass variance of one light shell, stabilizes the average mode.

The resulting branch supplies an analytic mass envelope, an exact captured rest-energy fraction, a derived late-time equation of state, a dark-density locking law, and a one-parameter small-scale cutoff family. The scale-covariant coupling itself is a model extension. It is not claimed to follow uniquely from the primitive projector axioms.

9.2 Additive one-shell obstruction

Let

$$\rho_V = \frac{\mu^4}{4\lambda}, \qquad m = \alpha\mu, \qquad \beta = \frac{\nu}{\mu^2}, \qquad r_D = \frac{\rho_{D0}}{\rho_V}. \tag{9.3}$$

Assume one shell of degeneracy g is populated at scale factor $a_\star$ by the energy-projected Gaussian

$$f(q) = \eta \exp\left[-\frac{\pi m^2}{\nu} - \frac{\pi q^2}{a_\star^2 \nu}\right], \qquad 0 < \eta \le 1. \tag{9.4}$$

In the nonrelativistic present limit,

$$r_D = \frac{\eta g \lambda \alpha \beta^{3/2}}{2\pi^3} e^{-\pi\alpha^2/\beta} a_\star^3, \tag{9.5}$$

and the present nonrelativistic momentum ratio is

$$u_0^2 := \frac{\langle p^2 \rangle_0}{m^2} = \frac{3\beta a_\star^2}{2\pi\alpha^2}. \tag{9.6}$$

Eliminating $\beta^{3/2} a_\star^3$ gives

$$u_0^3 = \frac{(3\pi)^{3/2}}{\sqrt{2}} \frac{r_D}{\eta g \lambda \alpha^4} e^{\pi\alpha^2/\beta}. \tag{9.7}$$

For an ultrarelativistic linear crossing, the unit-normalized produced energy is

$$\widetilde{\rho}_\star = \frac{g\nu^2}{4\pi^4}, \tag{9.8}$$

while the incoming kinetic energy is

$$K_\star = \frac{\nu^2}{6\lambda}. \tag{9.9}$$

Energy matching therefore implies

$$\eta g \lambda \le \frac{2\pi^4}{3}. \tag{9.10}$$

The additive locking condition is

$$-\mu^2 + 3\lambda \frac{\rho_{D0}}{m^2} > 0 \quad \Longrightarrow \quad \alpha^2 < \frac{3 r_D}{4}. \tag{9.11}$$

Combining (9.7)–(9.11) yields the conditional lower bound

$$\boxed{u_0^3 > \frac{4\sqrt{6}}{\pi^{5/2} r_D}.} \tag{9.12}$$

For

$$r_D = 0.3783102144, \tag{9.13}$$

this gives

$$\boxed{u_0 > 1.1397360601.} \tag{9.14}$$

The assumed nonrelativistic branch is therefore inconsistent. This is not a no-go theorem for every additive multi-level state. It rejects the specific one-shell, one-crossing, energy-matched mechanism that was being used to carry both abundance and locking.

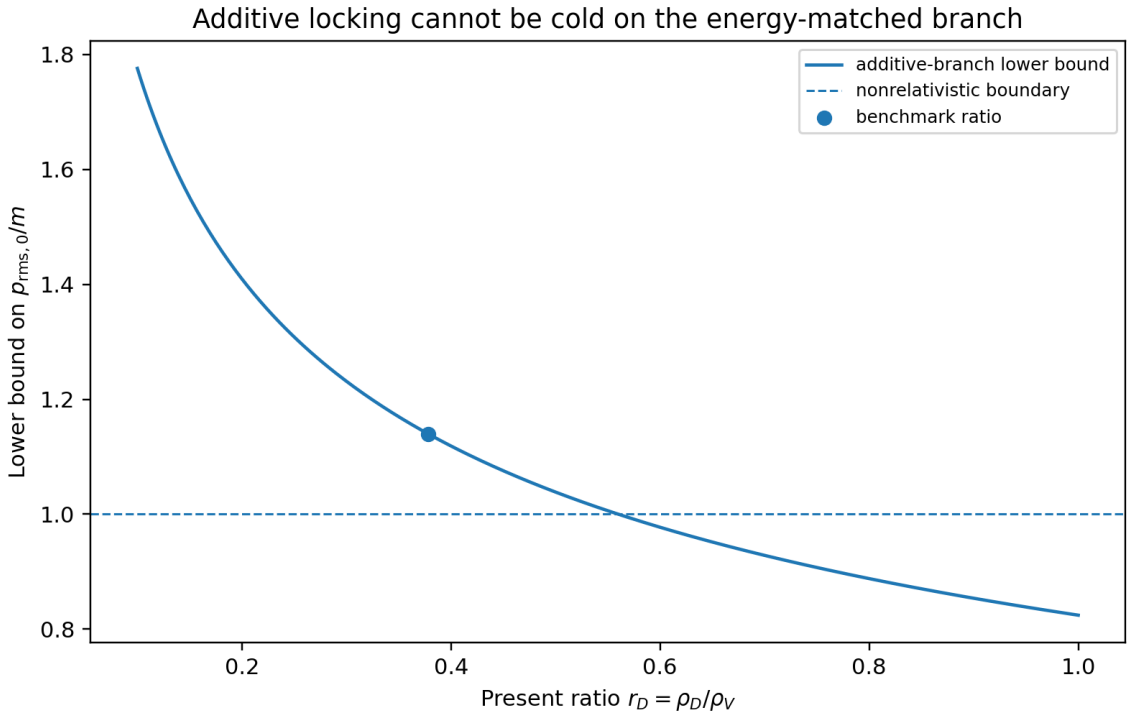


Figure 9.1. The abundance, energy-matching, and locking conditions force the additive one-shell ratio p_{rms}/m above unity at the declared present density ratio, contradicting the assumed cold branch.

9.3 Scale-covariant fiber action

Let

$$U(\phi) = \frac{\lambda}{4}(\phi^2 - v^2)^2, \qquad \mu^2 = \lambda v^2, \tag{9.15}$$

and let

$$\mathcal{L}_{\sigma,p} = \Lambda^2 \sum_{h\geq 1} p^{2\sigma h} \mathsf{D}_h. \tag{9.16}$$

The scale-covariant branch is defined by

$$\boxed{\begin{aligned} S_{\mathrm{SC}} = -\int \sqrt{-g}\, d^4x \Bigg[&\frac{1}{2}(\partial\phi)^2 + U(\phi) + \frac{1}{2}\langle \partial\chi, \partial\chi\rangle_p \\ &+ \frac{1}{2}\left\langle \chi, \left(m_0^2 - \mu^2 + \left(1 + \gamma\frac{\phi^2}{v^2}\right)\mathcal{L}_{\sigma,p}\right)\chi\right\rangle_p \Bigg]. \end{aligned}} \tag{9.17}$$

The level masses at the locked origin are

$$m_h^2 = m_0^2 - \mu^2 + \Lambda^2 p^{2\sigma h}, \tag{9.18}$$

with multiplicities

$$d_h = (p-1)p^{h-1}. \tag{9.19}$$

The average mode now changes the internal refinement scale itself. Its homogeneous equation is

$$\ddot{\phi} + 3H\dot{\phi} + \lambda\phi(\phi^2 - v^2) + \frac{\gamma\phi}{v^2}Y_D = 0, \tag{9.20}$$

where

$$Y_D := \sum_h \Lambda^2 p^{2\sigma h} X_h, \qquad X_h := \langle \chi_h^2 \rangle. \tag{9.21}$$

The origin curvature is

$$\boxed{M_\phi^2 = -\mu^2 + \frac{\gamma}{v^2}Y_D.} \tag{9.22}$$

For nonrelativistic levels with $m_h^2 \gg \mu^2, m_0^2$,

$$\Lambda^2 p^{2\sigma h} X_h = \rho_h \left[1 + O(w_h) + O\left(\frac{\mu^2 + m_0^2}{m_h^2}\right)\right]. \tag{9.23}$$

Thus the total cold dark density, including heavy levels, supplies the locking source.

9.4 Weighted nonadiabatic production

Near a crossing,

$$\phi(t) = \dot{\phi}_\star(t - t_\star) + O((t - t_\star)^2), \tag{9.24}$$

so that

$$\omega_{h,k}^2 = k^2 + m_h^2 + \frac{\gamma\Lambda^2 p^{2\sigma h}}{v^2}\dot{\phi}_\star^2 (t - t_\star)^2. \tag{9.25}$$

Define

$$A_\star := \frac{\sqrt{\gamma}|\dot{\phi}_\star|}{v}, \qquad \nu_h := A_\star \Lambda p^{\sigma h}. \tag{9.26}$$

The rapid-crossing spectrum is

$$\boxed{f_h(k) = \eta \exp\left[-\frac{\pi(k^2 + m_h^2)}{\nu_h}\right],} \tag{9.27}$$

where η is fixed by energy projection when the linearized spectrum would exceed the available kinetic energy:

$$\eta = \min\left(1, \frac{K_\star}{\widetilde{\rho}_\star}\right). \tag{9.28}$$

This is not an independent abundance fit. It is the shape-preserving rapid-crossing approximation with exact total-energy matching.

At high level,

$$m_h \simeq \Lambda p^{\sigma h}, \qquad \nu_h \simeq A_\star m_h. \tag{9.29}$$

The suppression exponent becomes linear rather than quadratic in mass:

$$\frac{\pi m_h^2}{\nu_h} \simeq \frac{\pi m_h}{A_\star}. \tag{9.30}$$

At the same time,

$$\frac{p_{\mathrm{rms},\star}}{m_h} \simeq \sqrt{\frac{3A_\star}{2\pi m_h}}, \tag{9.31}$$

so heavier populated levels are born colder.

9.5 Exact $\sigma = 1$ continuum theorem

For $\sigma = 1$, the endpoint law

$$N(m) = \frac{m}{\Lambda} - 1 \tag{9.32}$$

gives the continuum density of levels

$$\frac{dN}{dm} = \frac{1}{\Lambda}. \tag{9.33}$$

Using $\nu = Am$ and the occupation (9.27), the number and captured rest-density spectra are

$$\frac{dn}{dm} = \eta \frac{A^{3/2}}{8\pi^3\Lambda} m^{3/2} e^{-\pi m/A}, \tag{9.34}$$

$$\boxed{\frac{d\rho_R}{dm} = \eta \frac{A^{3/2}}{8\pi^3\Lambda} m^{5/2} e^{-\pi m/A}.} \tag{9.35}$$

The selected mass band peaks at

$$\boxed{m_{\mathrm{peak}} = \frac{5A}{2\pi}.} \tag{9.36}$$

The exact continuum integrals are

$$n_\star = C_N \eta \frac{A^4}{\Lambda}, \qquad C_N = \frac{3}{32\pi^5}, \tag{9.37}$$

$$\rho_{R,\star} = C_R \eta \frac{A^5}{\Lambda}, \qquad C_R = \frac{15}{64\pi^6}, \tag{9.38}$$

$$\rho_\star = C_E \eta \frac{A^5}{\Lambda}, \qquad C_E = \frac{32}{35\pi^7}, \tag{9.39}$$

$$p_\star = C_P \eta \frac{A^5}{\Lambda}, \qquad C_P = \frac{32}{315\pi^7}. \tag{9.40}$$

For example, the total-energy coefficient follows from the Tricomi representation and the beta integral:

$$\begin{aligned} C_E &= \frac{1}{2\pi^2} \int_0^\infty y^2 e^{-\pi y} I_\rho(\sqrt{y})\, dy \\ &= \frac{6}{\pi^7} B\left(\frac{3}{2}, 3\right) = \frac{32}{35\pi^7}. \end{aligned} \tag{9.41}$$

Consequently,

$$\boxed{w_\star = \frac{p_\star}{\rho_\star} = \frac{1}{9},} \tag{9.42}$$

$$\boxed{f_R := \frac{\rho_{R,\star}}{\rho_\star} = \frac{525\pi}{2048} = 0.8053399136\ldots,} \tag{9.43}$$

and

$$\frac{\rho_{R,\star}}{n_\star} = \frac{5A}{2\pi} = m_{\mathrm{peak}}. \tag{9.44}$$

A dyadic discrete tower with $\Lambda/A = 10^{-4}$ gives

$$f_R^{(2)} = 0.8054331018, \qquad w_\star^{(2)} = 0.1110616168, \tag{9.45}$$

and selects

$$\frac{m_{\mathrm{peak}}^{(2)}}{A} = 0.8192, \tag{9.46}$$

within one dyadic bin of (9.36). The absolute normalization retains the known log-periodic prime modulation, while the ratios are nearly envelope independent.

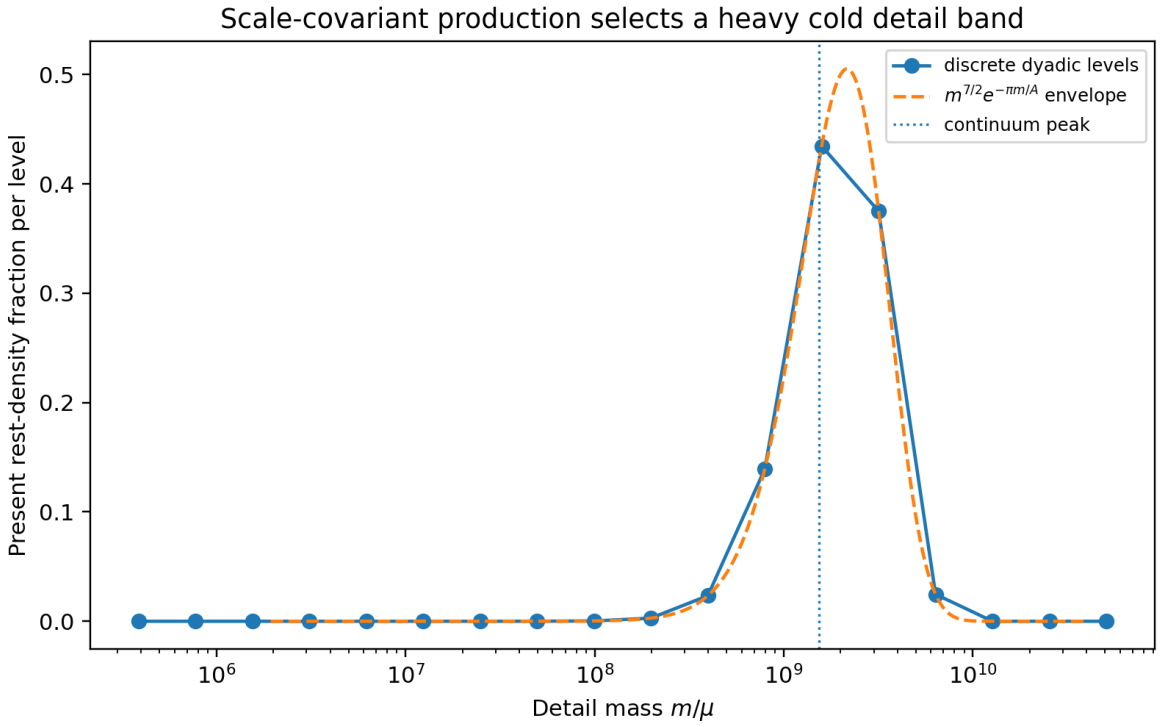


Figure 9.2. The discrete dyadic rest-density fractions follow the analytic scale-covariant mass envelope. The dominant mass is many orders of magnitude above the vacuum scale in the calibrated benchmark.

9.6 Abundance calibration and emergent hierarchy

On the captured branch,

$$\rho_{R,\star} = f_R K_\star. \tag{9.47}$$

After the produced band becomes nonrelativistic,

$$\rho_{D0} = f_R K_\star a_\star^3. \tag{9.48}$$

Define

$$F := \frac{K_\star}{\rho_V}. \tag{9.49}$$

The present abundance fixes

$$\boxed{F = \frac{r_D}{f_R a_\star^3}.} \tag{9.50}$$

The crossing scale and selected mass are then

$$\frac{A_\star}{\mu} = \sqrt{\frac{\gamma F}{2}}, \tag{9.51}$$

$$\boxed{\frac{m_{\rm peak}}{\mu} = \frac{5}{2\pi}\sqrt{\frac{\gamma r_D}{2 f_R a_\star^3}}.} \tag{9.52}$$

If the average field begins at rest at ϕ_i, then

$$F = \left(\frac{\phi_i}{v}\right)^4 - 2\left(\frac{\phi_i}{v}\right)^2, \tag{9.53}$$

so

$$\boxed{\frac{|\phi_i|}{v} = \sqrt{1+\sqrt{1+F}}.} \tag{9.54}$$

Thus one early production scale determines both the hierarchy between the vacuum and dark-matter masses and the required pre-crossing field amplitude.

For the benchmark

$$a_\star = 10^{-6}, \qquad \gamma = 16, \qquad p = 2, \qquad \frac{\Lambda}{A_\star} = 10^{-4}, \tag{9.55}$$

with (9.13), the discrete calculation gives

$$\frac{K_\star}{\rho_V} = 4.6969787254 \times 10^{17}, \tag{9.56}$$

$$\frac{A_\star}{\mu} = 1.9384486014 \times 10^9, \qquad \frac{|\phi_i|}{v} = 2.6179096175 \times 10^4, \tag{9.57}$$

$$\boxed{\frac{m^{(2)}_{\rm peak}}{\mu} = 1.5879770943 \times 10^9.} \tag{9.58}$$

The shape-preserving energy projection is

$$\eta = 3.8062069178 \times 10^{-21}. \tag{9.59}$$

This small occupation per microscopic state is compensated by the large refinement multiplicities. It is determined by energy conservation after the rapid-crossing spectrum is assembled.

9.7 Coldness theorem and cutoff scale

At $a \gg a_\star$, the scale-covariant continuum band obeys

$$\boxed{w_D(a) = \frac{a_\star^2}{5a^2},} \tag{9.60}$$

$$\boxed{c_{a,D}^2(a) = \frac{a_\star^2}{3a^2},} \tag{9.61}$$

and

$$\boxed{v_{\rm rms}(a) = \sqrt{\frac{3}{5}}\frac{a_\star}{a}.} \tag{9.62}$$

Unlike the additive one-shell result, the present speed decreases directly with the production scale. The reason is that the selected mass grows with the crossing rate while its fractional production momentum decreases as $m^{-1/2}$.

For $a_\star = 10^{-6}$,

$$w_{D0} = 1.9987378335 \times 10^{-13}, \qquad v_{\rm rms,0} = 7.7459666924 \times 10^{-7}. \tag{9.63}$$

The fluid Jeans estimate is

$$\frac{k_J(a)}{aH(a)} = \sqrt{\frac{3\Omega_m(a)}{2c_{a,D}^2(a)}}, \tag{9.64}$$

which gives

$$\boxed{\frac{k_J(a)}{H_0} = \frac{3}{\sqrt{2}}\frac{\sqrt{\Omega_{m0}a}}{a_\star}.} \tag{9.65}$$

At the present epoch,

$$\frac{k_{J0}}{H_0} = 1.1811011811 \times 10^6. \tag{9.66}$$

9.8 Density locking and transient acceleration

Using (9.23), the locked cold branch satisfies

$$\frac{M_\phi^2}{\mu^2} = -1 + \frac{\gamma}{4}r_D(a) + O(w_D) + O\left(\frac{\mu^2}{m_{\rm eff}^2}\right). \tag{9.67}$$

Hence present locking requires

$$\boxed{\gamma > \frac{4}{r_{D0}}.} \tag{9.68}$$

For the declared ratio, the right side is $10.573\ldots$. The benchmark choice $\gamma = 16$ gives

$$\boxed{\frac{M_\phi^2(1)}{\mu^2} = 0.5132408576.} \tag{9.69}$$

The exact discrete kinetic moments reconstruct

$$\frac{Y_D(1)}{\rho_V} = 0.3783102144 \tag{9.70}$$

within 2.3×10^{-13} of the declared dark-density ratio.

The cold asymptotic release scale is

$$a_{\mathrm{release}} = \left(\frac{\gamma r_{D0}}{4}\right)^{1/3}. \tag{9.71}$$

The exact discrete result is

$$\boxed{a_{\mathrm{release}} = 1.1480726020.} \tag{9.72}$$

For baryon ratio

$$r_{b0} = 0.0710237567, \tag{9.73}$$

the acceleration onset is

$$\boxed{a_{\mathrm{acc}} = 0.6081134084,} \tag{9.74}$$

with a cold asymptotic estimate differing by 1.94×10^{-4}.

At the locked point, the quadratic average-mode action remains

$$S_{\delta\phi}^{(2)} = \frac{1}{2}\int a^3 dt\, d^3x \left[\delta\dot{\phi}^2 - \frac{(\nabla\delta\phi)^2}{a^2} - M_\phi^2\delta\phi^2\right]. \tag{9.75}$$

Therefore

$$Q_{\delta\phi} = 1, \qquad c_{s,\phi}^2 = 1, \qquad M_\phi^2 > 0 \tag{9.76}$$

throughout the locked interval. Because every mass deformation is even in ϕ, the detail stress tensor also has no term linear in $\delta\phi$ at $\phi = 0$.

9.9 Linear growth and lensing benchmark

For $a \gg a_\star$, use $N = \log a$ and the controlled cold-fluid closure

$$\delta_b'' + \left(2 + \frac{d\log H}{dN}\right)\delta_b' = \frac{3}{2}S, \tag{9.77}$$

$$\delta_D'' + \left(2 + \frac{d\log H}{dN}\right)\delta_D' + \frac{c_{a,D}^2 k^2}{a^2H^2}\delta_D = \frac{3}{2}S, \tag{9.78}$$

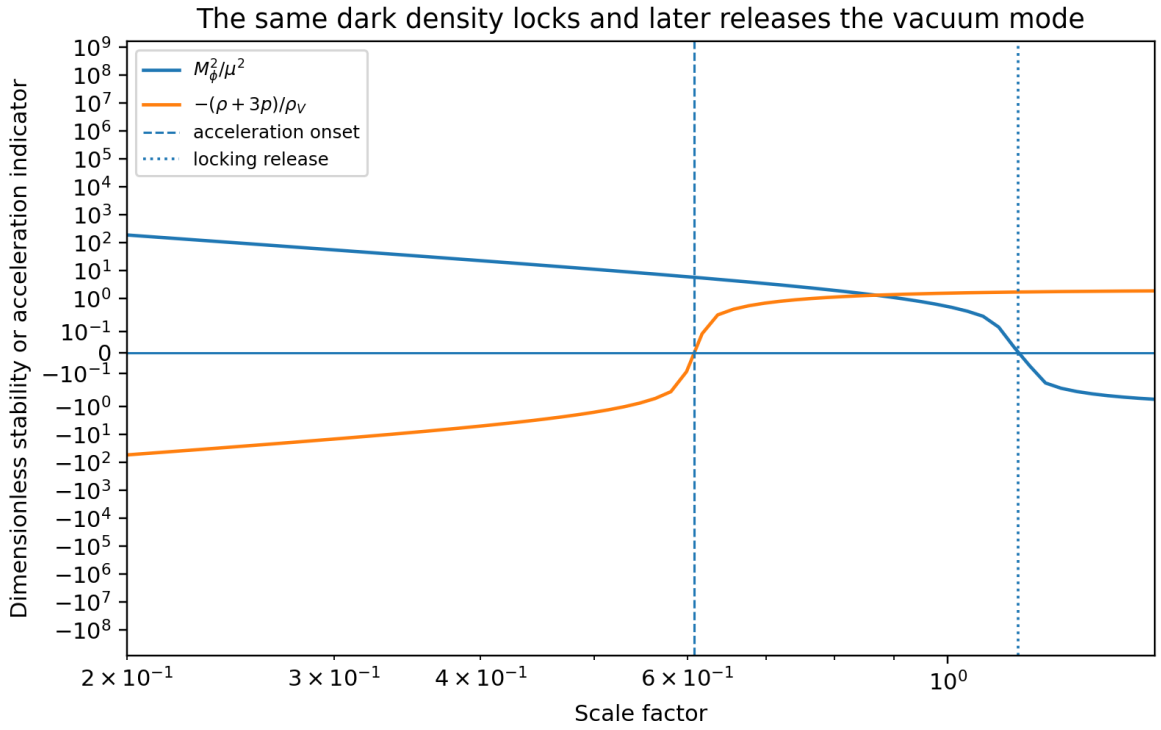


Figure 9.3. The weighted detail density stabilizes the average mode well before acceleration begins. Dilution eventually removes the locking curvature and releases the vacuum mode.

where

$$S = \Omega_b(a)\delta_b + \Omega_D(a)\delta_D. \tag{9.79}$$

The same source controls the subhorizon lensing potential to the accuracy of this closure.

Starting at $a_i = 10^{-3}$ with adiabatic growing data, the benchmark gives

k/H_0	$\delta_D/\delta_{D,\mathrm{cold}}$	lensing source / cold source
30	0.9999996989	0.9999997007
300	0.9999698901	0.9999700727
3000	0.9969926202	0.9970108463
10000	0.9669918115	0.9671896257
30000	0.7331608409	0.7346038042
100000	0.0159173474	0.0169297429

The first half-mode of the fluid lensing-source transfer is

$$\boxed{\frac{k_{1/2}}{H_0} = 4.4076353413 \times 10^4.} \tag{9.80}$$

For the conditional conversion $H_0 = 70\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$,

$$k_{1/2} = 10.2916022620\,\mathrm{Mpc}^{-1}. \tag{9.81}$$

This is a fluid-closure benchmark, not yet a complete collisionless Einstein–Boltzmann likelihood.

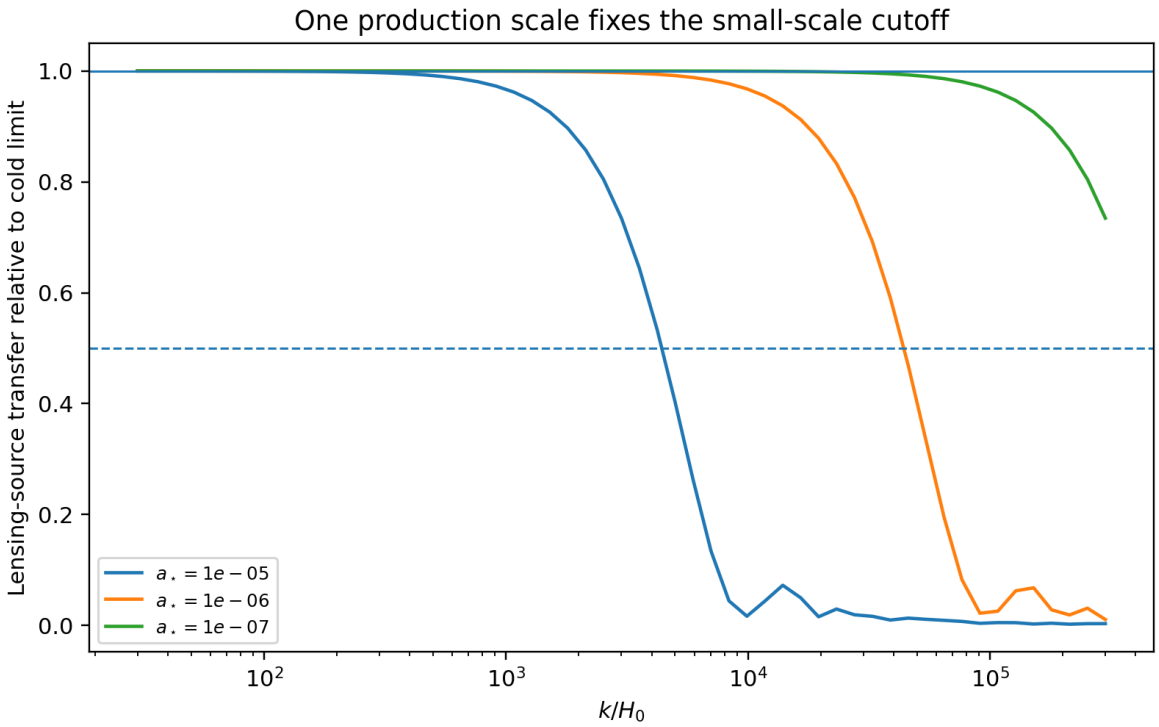


Figure 9.4. The production epoch controls the free-streaming cutoff without changing the late locking law. The high-k oscillations are those of the finite sound-speed fluid closure.

9.10 Verification ledger

The deterministic calculation records twenty passes and no failures. Selected checks are

statement	residual
additive abundance–velocity elimination identity	0
continuum C_E versus independent quadrature	1.62×10^{-12} relative
continuum C_P versus independent quadrature	3.49×10^{-11} relative
discrete versus continuum captured rest fraction	9.32×10^{-5}
discrete versus continuum production w	4.95×10^{-5}
present abundance reconstruction	0
exact versus asymptotic release scale	1.95×10^{-13}
exact versus asymptotic acceleration scale	1.93×10^{-4}
exact kinetic c_a^2 versus $a_\star^2/(3a^2)$	6.31×10^{-4} relative
large-scale lensing-source deviation at $k/H_0 = 300$	2.99×10^{-5}

The executable records are

$$R_{\mathrm{SC}} = \texttt{data/Scale_Covariant_Dark_Sector_Verification.json}, \tag{9.82}$$

$$C_{\mathrm{SC}} = \texttt{data/scale_covariant_dark_checks.csv}, \tag{9.83}$$

$$X_{\mathrm{SC}} = \texttt{verification/Scale_Covariant_Dark_Sector_Verification.py}. \tag{9.84}$$

9.11 Closure ledger and remaining burden

statement	status	boundary
additive one-shell cold-lock obstruction	proved conditional	rapid, energy-matched, dominant-shell branch
scale-covariant fiber action	defined	one new coupling γ
weighted production spectrum	derived	linear rapid crossing
$\sigma = 1$ continuum mass theorem	proved	high-level density-of-states envelope
energy-projected rest fraction	proved	shape-preserving capture approximation
late w_D, c_a^2, and $v_{\rm rms}$ laws	proved asymptotically	collisionless nonrelativistic regime
density locking and release	closed controlled	heavy-level limit and regular locked phase
linear growth and lensing cutoff	numerical controlled	two-fluid sound-speed closure
full collisionless hierarchy	open	momentum and anisotropic-stress multipoles
radiative stability of scale covariance	open	quantum effective action
microscopic derivation of γ and $a_\star$	open	gravitational refinement dynamics
observational likelihood and priority audit	open	independent implementation and data comparison

The principal result is not that every earlier dark-sector formula survives unchanged. The hard consistency audit rejects the additive dominant-shell mechanism and identifies a minimal viable replacement:

$$\boxed{\begin{array}{c}\text{average-mode dilation of the refinement spectrum}\\ \Longrightarrow\ \text{heavy cold detail band}\\ \Longrightarrow\ \text{density-weighted vacuum locking}\end{array}} \tag{9.85}$$

A complete physical treatment therefore requires the full collisionless Einstein–Boltzmann hierarchy for the distribution (9.27), with one parameter set propagated through recombination, matter power, weak lensing, and the transient-release background without sector-by-sector renormalization.

CHAPTER 10

Scale-Covariant Collisionless Boltzmann Dynamics and Late-Time Observables

From a distribution to observables. The third part carries one prescribed distribution through momentum-resolved transport, photon and neutrino hierarchies, recombination, reionization, line-of-sight projection, and CMB lensing. These are controlled templates, not a precision replacement for production codes.

10.1 Scope and objective

The scale-covariant refinement dark sector fixes a discrete mass spectrum, a production distribution, a cold late-time limit, and a density-dependent locking source. The preceding growth calculation replaced that distribution by one adiabatic sound speed. That fluid closure is adequate on large scales but does not retain momentum dependence, free streaming, or the higher angular moments of the collisionless distribution.

This chapter develops the kinetic layer in two stages. First, it derives the complete linear collisionless hierarchy for every retained refinement level and momentum node. Second, it specializes that hierarchy to a controlled post-recombination, subhorizon Einstein–Vlasov–Poisson system and propagates one fixed scale-covariant distribution into matter-power, redshift-space-growth, weak-lensing, and CMB-lensing templates.

The numerical section begins at

$$a_i = 10^{-3}. \tag{10.1}$$

Adiabatic perturbations are prescribed at that surface. The calculation therefore measures the *incremental post-recombination transfer* generated between a_i and today; it does not evolve free streaming from the production event to a_i. Radiation is retained in the background expansion but its perturbations are not included in the late-time Poisson source. Consequently the section does not compute primary CMB temperature or polarization spectra, recombination kinetics, neutrino transport, or superhorizon metric evolution. The CMB result below is a late-time lensing template and a scale-separation diagnostic, not a primary-anisotropy likelihood.

10.2 Scale-covariant phase-space distribution

For each retained detail level h, let

$$m_h = \Lambda p^{\sigma h}, \qquad d_h = (p-1)p^{h-1}, \tag{10.2}$$

within the high-level scale-covariant branch. The rapid-crossing distribution has the form

$$\bar{f}_h(q) = \eta \exp\left[-\frac{\pi m_h^2}{\nu_h} - \frac{\pi q^2}{a_\star^2 \nu_h}\right], \qquad \nu_h = A_\star m_h. \tag{10.3}$$

For the benchmark used here, $A_\star$ is the mass unit, $p = 2$, $\sigma = 1$, and

$$\frac{\Lambda}{A_\star} = 10^{-4}, \qquad a_\star = 10^{-6}. \tag{10.4}$$

The common energy-projection normalization η cancels from every linear transfer ratio.

The conformal energy is

$$\epsilon_h(q,a) = \sqrt{q^2 + a^2 m_h^2}, \tag{10.5}$$

and the background density and pressure are

$$\rho_h(a) = \frac{d_h}{2\pi^2 a^4} \int_0^\infty q^2 \epsilon_h \bar{f}_h \, dq, \tag{10.6}$$

$$p_h(a) = \frac{d_h}{6\pi^2 a^4} \int_0^\infty \frac{q^4}{\epsilon_h} \bar{f}_h \, dq. \tag{10.7}$$

The numerical phase grid retains the ten levels carrying all but less than 10^{-6} of the present rest-density weight. Their retained weights sum to unity to 2.3×10^{-16}.

10.3 Complete linear collisionless hierarchy

Use conformal Newtonian gauge,

$$ds^2 = a^2(\eta)\left[-(1+2\Psi)d\eta^2 + (1-2\Phi)d\mathbf{x}^2\right]. \tag{10.8}$$

Write the perturbed distribution as

$$f_h(\mathbf{k}, q, \mu, \eta) = \bar{f}_h(q)\left[1 + \mathcal{F}_h(\mathbf{k}, q, \mu, \eta)\right], \tag{10.9}$$

where $\mu = \widehat{\mathbf{k}} \cdot \widehat{\mathbf{q}}$, and expand

$$\mathcal{F}_h = \sum_{\ell \geq 0} (-i)^\ell (2\ell + 1) F_{h\ell}(k, q, \eta) P_\ell(\mu). \tag{10.10}$$

Set

$$D_h(q) := \frac{d \log \bar{f}_h}{d \log q}. \tag{10.11}$$

The linear collisionless Boltzmann equation gives

$$\dot{F}_{h0} = -k \frac{q}{\epsilon_h} F_{h1} + \dot{\Phi} D_h, \tag{10.12}$$

$$\dot{F}_{h1} = \frac{kq}{3\epsilon_h}\left(F_{h0} - 2F_{h2}\right) - \frac{k\epsilon_h}{3q} \Psi D_h, \tag{10.13}$$

$$\dot{F}_{h\ell} = \frac{kq}{(2\ell+1)\epsilon_h}\left[\ell F_{h,\ell-1} - (\ell+1)F_{h,\ell+1}\right], \qquad \ell \geq 2. \tag{10.14}$$

These equations retain the complete angular hierarchy of the produced distribution. No sound speed is inserted.

The stress-energy moments are

$$\delta\rho_h = \frac{d_h}{2\pi^2 a^4}\int q^2\epsilon_h \bar{f}_h F_{h0}\, dq, \tag{10.15}$$

$$\delta p_h = \frac{d_h}{6\pi^2 a^4}\int \frac{q^4}{\epsilon_h}\bar{f}_h F_{h0}\, dq, \tag{10.16}$$

$$(\rho_h + p_h)\theta_h = \frac{d_h k}{2\pi^2 a^4}\int q^3 \bar{f}_h F_{h1}\, dq, \tag{10.17}$$

$$(\rho_h + p_h)\sigma_h = \frac{d_h}{3\pi^2 a^4}\int \frac{q^4}{\epsilon_h}\bar{f}_h F_{h2}\, dq. \tag{10.18}$$

Thus density, pressure, velocity, and anisotropic stress are all projections of one hierarchy.

10.4 Post-recombination Vlasov–Poisson specialization

For the controlled late-time calculation, the time derivative of the spatial metric potential is neglected relative to the subhorizon force term. The hierarchy becomes a Vlasov system sourced by

$$k^2\Psi = -4\pi G a^2\left(\rho_b\delta_b + \sum_h \delta\rho_h\right). \tag{10.19}$$

The baryon equations are

$$\dot{\delta}_b = -\theta_b, \qquad \dot{\theta}_b + \mathcal{H}\theta_b = k^2\Psi. \tag{10.20}$$

Let $N = \log a$, $E(a) = H(a)/H_0$, and

$$\kappa(a) := \frac{k}{aH} = \frac{k/H_0}{aE(a)}. \tag{10.21}$$

Then the implemented hierarchy is

$$\frac{dF_{h0}}{dN} = -\kappa\frac{q}{\epsilon_h}F_{h1}, \tag{10.22}$$

$$\frac{dF_{h1}}{dN} = \kappa\left[\frac{q}{3\epsilon_h}(F_{h0} - 2F_{h2}) - \frac{\epsilon_h}{3q}\Psi D_h\right], \tag{10.23}$$

with the corresponding N-form of (10.14).

The background is not replaced by exact dust. It uses the same phase-space quadrature as the hierarchy:

$$E(a)^2 = \Omega_{r0}a^{-4} + \Omega_{b0}a^{-3} + \Omega_{D0}\frac{\rho_D(a)}{\rho_D(1)} + \Omega_{V0}. \tag{10.24}$$

The locked average mode has no linear stress perturbation in this phase and enters only through Ω_{V0}.

10.5 Adiabatic growing initialization

At $a_i = 10^{-3}$, the retained tower is already highly nonrelativistic. Set

$$F_{h0}(q, N_i) = \delta_i, \qquad \delta_i = a_i. \tag{10.25}$$

Let

$$u_i := \frac{\theta_i}{\mathcal{H}_i} = -\delta_i, \qquad \kappa_i := \frac{k}{\mathcal{H}_i}. \tag{10.26}$$

The momentum-dependent boost initialization is

$$F_{h1}(q, N_i) = \frac{u_i}{\kappa_i}\left[-\frac{D_h(q)}{3v_h(q, N_i)}\right], \qquad v_h := \frac{q}{\epsilon_h}. \tag{10.27}$$

For the Gaussian production distribution,

$$D_h = -2\pi x^2, \qquad x = \frac{q}{a_\star\sqrt{\nu_h}}, \tag{10.28}$$

and

$$\frac{\int_0^\infty x^2 e^{-\pi x^2}(-D_h/3)\,dx}{\int_0^\infty x^2 e^{-\pi x^2}\,dx} = 1. \tag{10.29}$$

Hence the hierarchy reproduces the growing-mode continuity normalization in the cold limit without assigning one artificial velocity to every momentum node.

10.6 Multipole and momentum closure

Momentum integration uses generalized Gauss–Laguerre quadrature with weight $y^{1/2}e^{-y}$. The declared production transfer is converged to

$$\epsilon_q = 3.05 \times 10^{-5} \tag{10.30}$$

when increasing the quadrature from eight to ten nodes per retained level.

For the angular tail, define the free-streaming distance

$$s_h(q, N) := \int_{N_i}^{N} \frac{q/\epsilon_h}{aE(a)}\,dN. \tag{10.31}$$

When $ks_h > L + 1$, use the spherical-Bessel recurrence closure

$$F_{h,L+1} = \frac{2L+1}{ks_h}F_{hL} - F_{h,L-1}. \tag{10.32}$$

The free hierarchy reproduces $F_\ell = j_\ell(x)$ through the tested low multipoles with maximum residual

$$\epsilon_{\mathrm{fs}} = 9.52 \times 10^{-9}. \tag{10.33}$$

The $L = 6$ to $L = 8$ present-transfer difference on the declared convergence sample is

$$\epsilon_L = 3.62 \times 10^{-5}. \tag{10.34}$$

10.7 Matter and Weyl transfer functions

Define the total matter contrast and the Poisson/Weyl source by

$$\delta_m = \frac{\rho_b \delta_b + \rho_D \delta_D}{\rho_b + \rho_D}, \tag{10.35}$$

$$\begin{aligned} S_W = {} & \Omega_{b0} a^{-1} \delta_b + \Omega_{D0} a^2 \frac{\rho_D(a)}{\rho_D(1)} \delta_D \\ & - \frac{3}{2} \Omega_{D0} a^2 \frac{\rho_D(a)}{\rho_D(1)} (1 + w_D) \sigma_D. \end{aligned} \tag{10.36}$$

The last term is the collisionless gravitational-slip contribution. It is retained in the lensing source even though it is numerically tiny for the declared cold benchmark. Relative to a pressureless reference evolved from the same initial amplitude, set

$$T_m(k,a) = \frac{\delta_m(k,a)}{\delta_m^{\rm cold}(a)}, \qquad T_W(k,a) = \frac{S_W(k,a)}{S_W^{\rm cold}(a)}. \tag{10.37}$$

The linear matter-power ratio is

$$\frac{P_m(k,a)}{P_m^{\rm cold}(k,a)} = T_m(k,a)^2. \tag{10.38}$$

The post-recombination transfer is unity on large scales and develops a finite incremental free-streaming cutoff. The amplitude half-mode is

$$\boxed{\frac{k_{1/2}^{(T)}}{H_0} = 4.71448 \times 10^4,} \tag{10.39}$$

while the power half-mode is

$$\boxed{\frac{k_{1/2}^{(P)}}{H_0} = 3.14535 \times 10^4.} \tag{10.40}$$

For the conditional conversion $H_0 = 70\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$,

$$\boxed{k_{1/2}^{(T)} = 11.0081\,\mathrm{Mpc}^{-1}, \qquad k_{1/2}^{(P)} = 7.34423\,\mathrm{Mpc}^{-1}.} \tag{10.41}$$

These conversions are benchmark conventions rather than a likelihood fit.

Selected present values are

k/H_0	T_m	$P_m/P_m^{\rm cold}$	$T_m^{\rm fluid}$
3.15×10^2	0.999960	0.999921	0.999967
2.72×10^3	0.997044	0.994096	0.997542
2.86×10^4	0.747869	0.559307	0.756672
5.14×10^4	0.451041	0.203438	0.380038
9.26×10^4	0.161800	0.026179	0.019458
3.00×10^5	0.021197	0.000449	0.010676

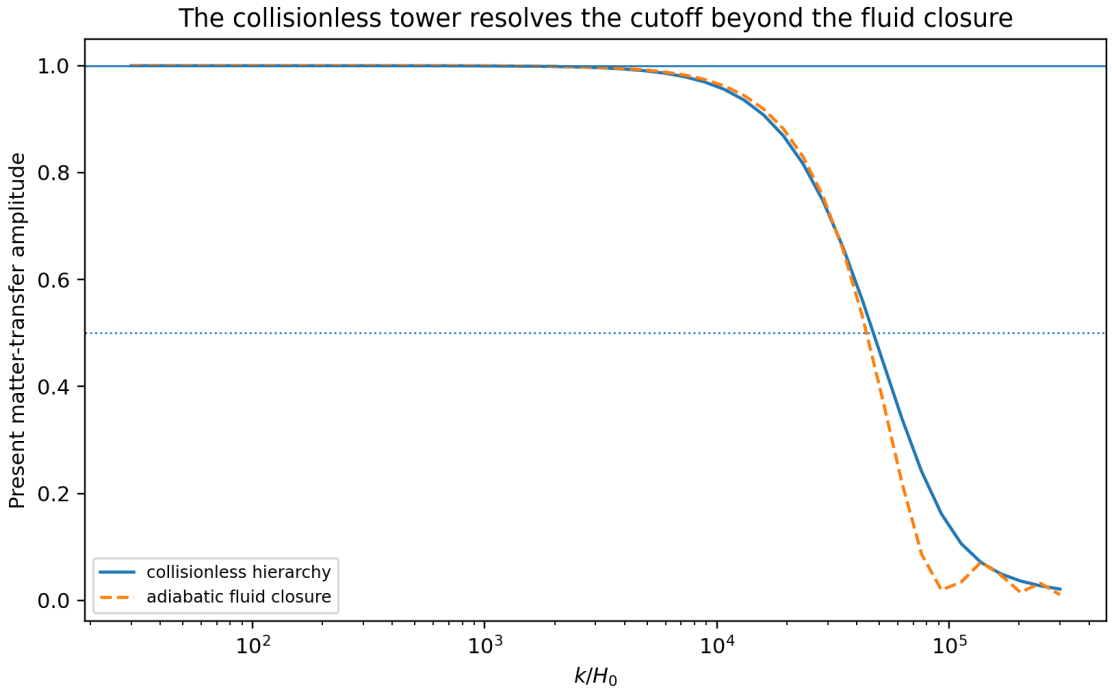


Figure 10.1. The collisionless hierarchy and the earlier adiabatic-fluid closure agree at large scale but differ materially after free streaming populates higher angular moments.

The hierarchy is not equivalent to replacing the distribution by a single sound speed. The difference changes sign across the cutoff because the collisionless response contains phase mixing rather than one coherent acoustic oscillation.

10.8 Redshift-space growth and anisotropic stress

On the common logarithmic scale-factor grid, define

$$f_m(k,a) = \frac{d\log \delta_m(k,a)}{d\log a}. \tag{10.42}$$

For equal primordial normalization, the scale-dependent redshift-space amplitude ratio is

$$\mathcal{R}_{f\sigma_8}(k,a) = \frac{f_m \delta_m}{f_m^{\rm cold}\delta_m^{\rm cold}}. \tag{10.43}$$

At $z = 0$, the benchmark gives

$$\mathcal{R}_{f\sigma_8} = 0.9970 \quad \text{at} \quad k/H_0 \simeq 2.72\times 10^3, \tag{10.44}$$

$$\mathcal{R}_{f\sigma_8} = 0.7471 \quad \text{at} \quad k/H_0 \simeq 2.86\times 10^4, \tag{10.45}$$

$$\mathcal{R}_{f\sigma_8} = 0.1603 \quad \text{at} \quad k/H_0 \simeq 9.26\times 10^4. \tag{10.46}$$

The computed late-time anisotropic stress remains finite and is below 1.5×10^{-12} on the selected observable table. It is dynamically retained even though it is numerically small for this very cold benchmark.

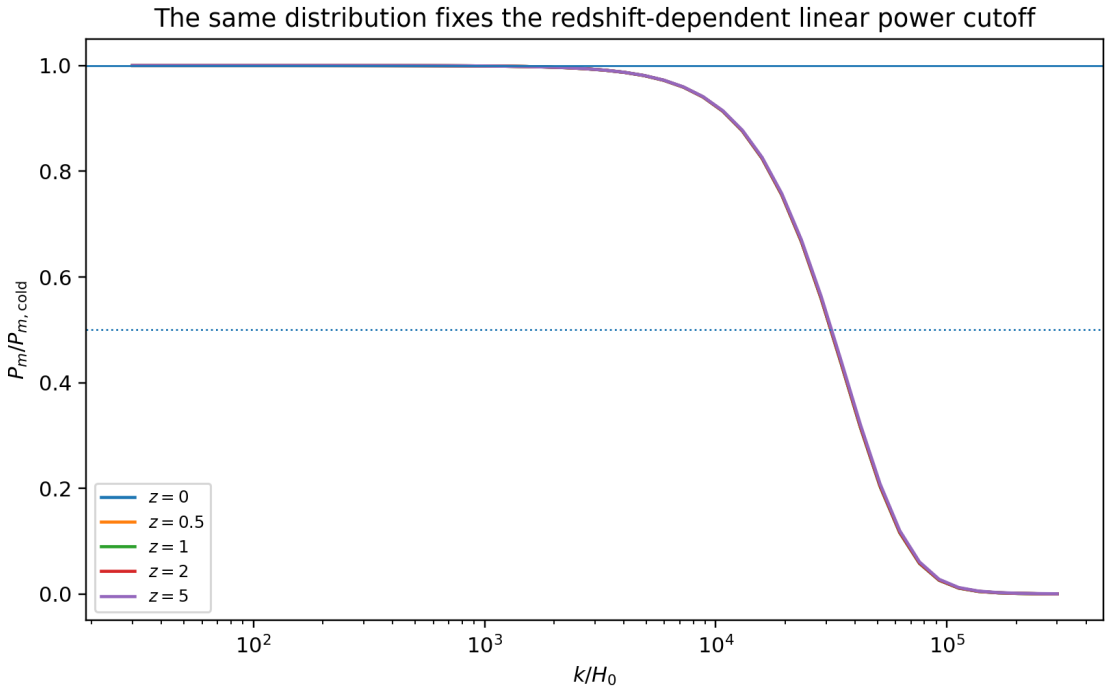


Figure 10.2. The power cutoff develops continuously after the declared initial epoch. The large scales remain locked to the cold reference while the high-k modes lose coherent growth.

10.9 Weak-lensing and CMB-lensing templates

For a flat source plane at comoving distance χ_s, define the controlled linear lensing ratio

$$\mathcal{R}_\ell^\kappa = \frac{\int_0^{\chi_s} d\chi\, W(\chi,\chi_s)^2 \chi^{-2} P_{\mathrm{cold}}(k,z) T_W(k,z)^2}{\int_0^{\chi_s} d\chi\, W(\chi,\chi_s)^2 \chi^{-2} P_{\mathrm{cold}}(k,z)}, \tag{10.47}$$

where

$$k = \frac{\ell + 1/2}{\chi}. \tag{10.48}$$

The reference weight uses the same cold growth and the declared positive equality filter

$$T_{\mathrm{eq}}(k) = \frac{1}{1 + (k/k_{\mathrm{eq}})^2}, \qquad \frac{k_{\mathrm{eq}}}{H_0} = \Omega_{m0}\sqrt{\frac{2}{\Omega_{r0}}}. \tag{10.49}$$

This is a normalized linear diagnostic. It is not a nonlinear shear likelihood and does not include survey masks, intrinsic alignments, nonlinear baryonic feedback, or post-Born corrections.

For a galaxy source at $z_s = 0.5$,

$$\mathcal{R}^\kappa_{1000} = 0.981331, \qquad \mathcal{R}^\kappa_{3000} = 0.863481, \qquad \mathcal{R}^\kappa_{5000} = 0.702128. \tag{10.50}$$

For $z_s = 1$,

$$\mathcal{R}^{\kappa}_{1000} = 0.993772, \qquad \mathcal{R}^{\kappa}_{3000} = 0.949092, \qquad \mathcal{R}^{\kappa}_{5000} = 0.874908. \tag{10.51}$$

For a CMB source plane at the declared initial redshift $z_s = 999$,

$$\mathcal{R}^{\kappa}_{1000} = 0.999615, \qquad \mathcal{R}^{\kappa}_{2000} = 0.998474, \qquad \mathcal{R}^{\kappa}_{3000} = 0.996591. \tag{10.52}$$

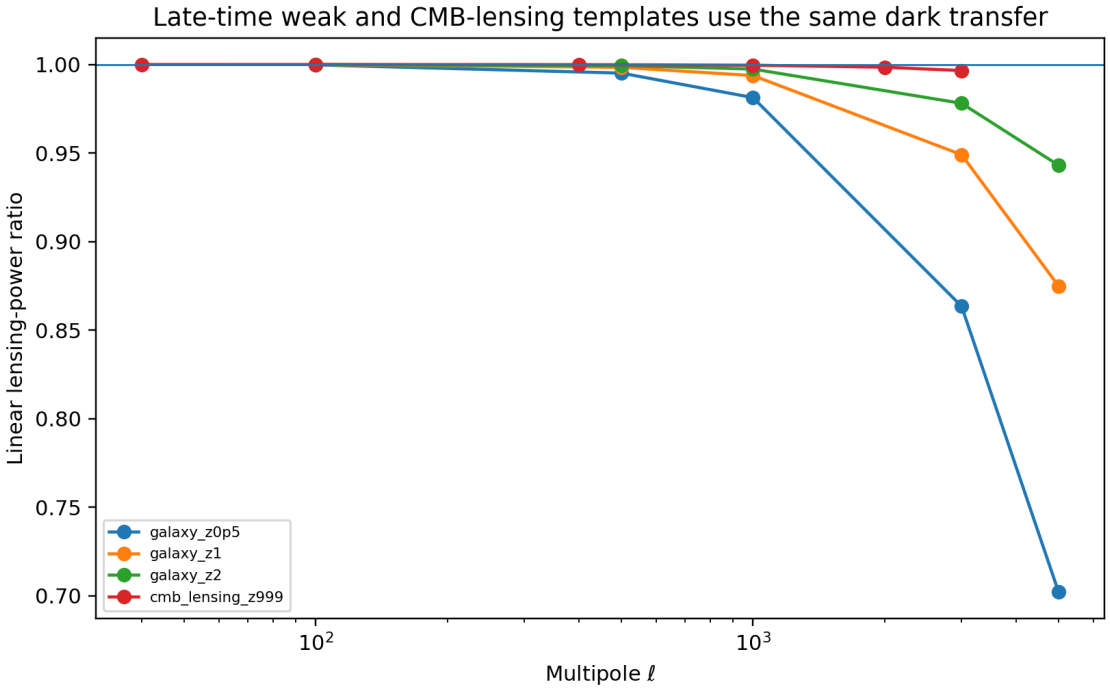


Figure 10.3. The late-time linear lensing response follows from the same collisionless Weyl transfer function without changing the dark parameters between clustering and lensing.

10.10 Primary-CMB scale separation

The comoving distance of the controlled background to $a = 10^{-3}$ maps $\ell = 3000$ to

$$\left.\frac{k}{H_0}\right|_{\ell=3000} = 957.699. \tag{10.53}$$

This is only

$$\boxed{\frac{k_{\ell=3000}}{k^{(P)}_{1/2}} = 0.0304481} \tag{10.54}$$

of the late-time power half-mode. Thus the declared cutoff is parametrically separated from the primary acoustic scale. This does not replace a primary TT, TE, or EE calculation. It shows only a scale separation inside the declared late-time model. Pre-a_i free streaming, photon and neutrino metric driving, and recombination physics must be evolved before any statement about primary anisotropies is justified.

10.11 Verification ledger

The deterministic implementation records

$$\boxed{\text{20 PASS,} \qquad \text{0 FAIL.}} \tag{10.55}$$

The principal numerical controls are

control	coarse difference	fine difference
momentum quadrature	1.41×10^{-4}	3.05×10^{-5}
multipole truncation	3.62×10^{-5}	3.62×10^{-5}
time-step refinement	5.41×10^{-8}	1.89×10^{-8}
cold-production limit	–	4.01×10^{-4}

The time-step sequence has an effective fitted order 1.52 at the reported precision because the remaining error is already dominated by the moving angular closure and phase-space truncation. The fine transfer difference itself is below 2×10^{-8}.

10.12 Closure ledger

statement	status	boundary
massive collisionless hierarchy for every retained level and momentum node	derived and implemented	linear perturbations
phase-space stress tensor including anisotropic stress	closed	declared quadrature
self-consistent baryon–tower Poisson source	closed controlled	quasi-static post-recombination regime
large-scale cold limit	numerically closed	declared hierarchy grid
post-recombination incremental transfer cutoff	numerically closed	prescribed data at a_i
scale-dependent $f\sigma_8$ template	closed controlled	linear theory
galaxy weak-lensing ratio	closed controlled	linear Limber template
CMB-lensing ratio	closed controlled	linear late-time lensing only
primary TT, TE, EE spectra	open	photon collisions and polarization
massless-neutrino hierarchy and recombination coupling	open	full Einstein–Boltzmann system
nonlinear matter power and halo formation	open	nonlinear solver or calibrated emulator
observational likelihood and parameter inference	open	external data and nuisance model

10.13 Synthesis and unresolved requirements

The scale-covariant refinement distribution now determines, without replacing it by one sound speed,

$$\boxed{\bar{f}_h(q) \Longrightarrow \{F_{h\ell}(q,k,a)\} \Longrightarrow \{T_m, T_W, f\sigma_8, \mathcal{R}^\kappa_\ell\}.} \tag{10.56}$$

The same mass spectrum and production scale control dark clustering, redshift-space growth, galaxy lensing, and CMB lensing.

The unresolved kinetic and observational requirements are:

$$\boxed{\begin{array}{c}\text{add photons, polarization, baryons, and massless neutrinos,}\\ \text{couple the hierarchy through recombination and reionization,}\\ \text{compute TT, TE, EE, lensing, and } P_m(k) \text{ together,}\\ \text{then perform one joint likelihood without retuning the dark sector.}\end{array}} \tag{10.57}$$

CHAPTER 11

Controlled Production-to-Primary-CMB Einstein–Boltzmann Dynamics

11.1 Scope and objective

The scale-covariant refinement dark sector now supplies a mass spectrum, a production distribution, a cold late-time limit, a collisionless perturbation hierarchy, and controlled late-time matter and lensing observables. The remaining major linear-cosmology gap was the interval between production and last scattering. In particular, the preceding late-time calculation did not jointly evolve photon collisions, polarization, baryons, massless neutrinos, metric driving, recombination, and the massive refinement hierarchy.

This chapter constructs that coupled system. It evolves one unchanged dark-sector parameter set from

$$a_\star = 10^{-6} \tag{11.1}$$

through recombination and beyond the visibility peak. The numerical state contains photon temperature and polarization multipoles, baryon density and velocity, massless-neutrino multipoles, and momentum-resolved massive refinement multipoles. The metric potentials are obtained from the Einstein momentum and anisotropic-stress equations, while the energy equation is retained as an independent diagnostic.

The resulting TT, TE, and EE spectra are *controlled recombination-era templates*. They demonstrate acoustic oscillations, polarization generation, neutrino and dark-sector metric feedback, finite-thickness damping, and a consistent geometric acoustic ruler. They are not precision CMB likelihood spectra. The calculation uses Saha hydrogen recombination, no reionization, no multilevel helium or hydrogen kinetics, no late integrated Sachs–Wolfe source, no lensing remapping, and a finite-width narrow-source projection rather than a production-quality line-of-sight integrator.

11.2 One unchanged background and refinement distribution

Use units in which wavenumbers and conformal rates are measured in H_0. The background is

$$E(a)^2 = \Omega_{\gamma 0} a^{-4} + \Omega_{\nu 0} a^{-4} + \Omega_{b0} a^{-3} + \Omega_{D0} \frac{\rho_D(a)}{\rho_D(1)} + \Omega_{V0}, \tag{11.2}$$

with

$$\frac{\mathcal{H}}{H_0} = aE(a). \tag{11.3}$$

The declared numerical fractions are

$$\Omega_{b0} = 0.049, \qquad \Omega_{D0} = 0.261, \qquad \Omega_{r0} = 9.0 \times 10^{-5}, \tag{11.4}$$

with $T_{\gamma 0} = 2.7255\,\mathrm{K}$ and $H_0 = 70\,\mathrm{km\,s^{-1}\,Mpc^{-1}}$ for the conditional conversion to physical units. The blackbody photon density is computed from $T_{\gamma 0}$,

$$\Omega_{\gamma 0} = 5.0468884259 \times 10^{-5}, \tag{11.5}$$

and the remaining declared radiation density is assigned to a massless collisionless-neutrino sector,

$$\Omega_{\nu 0} = 3.9531115741 \times 10^{-5}. \tag{11.6}$$

For the scale-covariant refinement tower,

$$m_h = A_\star \frac{\Lambda}{A_\star} p^{\sigma h}, \qquad d_h = (p-1)\,p^{h-1}, \tag{11.7}$$

with

$$p = 2, \qquad \sigma = 1, \qquad \frac{\Lambda}{A_\star} = 10^{-4}. \tag{11.8}$$

The produced distribution is proportional to

$$\bar{f}_h(q) \propto \exp\left[-\pi \frac{m_h}{A_\star} - \pi \frac{q^2}{a_\star^2 A_\star m_h}\right]. \tag{11.9}$$

Its common normalization is fixed by Ω_{D0} and cancels from the normalized perturbation transfer. The background moments are

$$\epsilon_h(q,a) = \sqrt{q^2 + a^2 m_h^2}, \tag{11.10}$$

$$\rho_D(a) = \frac{1}{2\pi^2 a^4} \sum_h d_h \int_0^\infty q^2 \epsilon_h \bar{f}_h \, dq, \tag{11.11}$$

$$p_D(a) = \frac{1}{6\pi^2 a^4} \sum_h d_h \int_0^\infty \frac{q^4}{\epsilon_h} \bar{f}_h \, dq. \tag{11.12}$$

Seven levels, $h = 9, \ldots, 15$, carry the retained present rest-density weight at the declared 10^{-4} threshold. Two generalized Gauss–Laguerre nodes per level give fourteen massive phase-space nodes in the main solve.

At the visibility peak, the tower equation of state is

$$\boxed{w_D(a_{\mathrm{vis}}) = 3.3085372607 \times 10^{-7}.} \tag{11.13}$$

The produced refinement sector is therefore already effectively cold during last scattering, while its momentum dependence and anisotropic stress remain explicitly present in the hierarchy.

11.3 Saha recombination, opacity, and visibility

Let

$$T_\gamma(a) = \frac{T_{\gamma 0}}{a}, \qquad n_H(a) = \frac{(1 - Y_p)\Omega_{b0}\rho_{c0}}{m_p a^3}, \tag{11.14}$$

with $Y_p = 0.24$. The Saha ionization fraction obeys

$$\frac{x_e^2}{1-x_e} = \frac{1}{n_H}\left(\frac{m_e k_B T_\gamma}{2\pi\hbar^2}\right)^{3/2} \exp\left(-\frac{\chi_H}{k_B T_\gamma}\right). \tag{11.15}$$

Define the positive Thomson interaction rate

$$\dot{\kappa} = a n_e \sigma_T c, \tag{11.16}$$

where the dot denotes $d/d\eta$. For $N = \log a$, the optical depth to the present is

$$\tau(N) = \int_N^0 \frac{\dot{\kappa}}{\mathcal{H}}\, dN, \tag{11.17}$$

and the visibility in logarithmic scale-factor measure is

$$g_N(N) = e^{-\tau}\frac{\dot{\kappa}}{\mathcal{H}}. \tag{11.18}$$

The numerical integral gives

$$\int g_N\, dN = 0.9999798979, \tag{11.19}$$

with peak

$$\boxed{z_{\rm vis} = 1286.261163, \qquad a_{\rm vis} = 7.7684313701 \times 10^{-4}.} \tag{11.20}$$

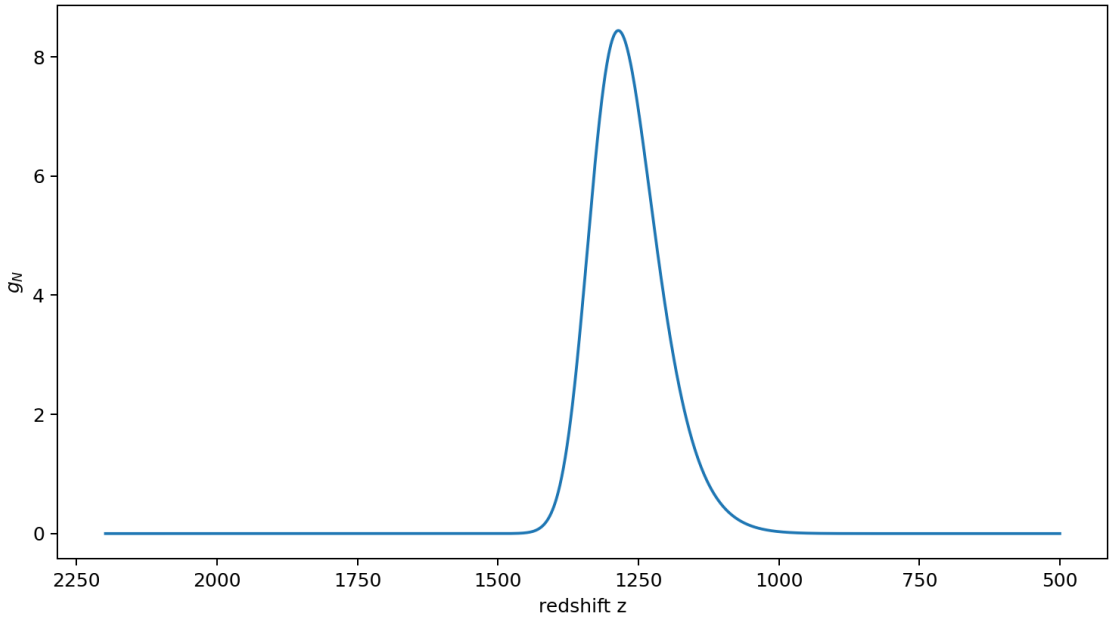


Figure 11.1. The controlled Saha visibility function in $d\log a$ measure. The construction omits multilevel recombination, helium corrections beyond the fixed hydrogen fraction, and reionization.

11.4 Metric conventions and Einstein closure

Use conformal Newtonian gauge,

$$ds^2 = a^2(\eta)\left[-(1+2\Psi)d\eta^2 + (1-2\Phi)d\mathbf{x}^2\right]. \tag{11.21}$$

Let

$$\mathcal{Q} = \sum_s (\rho_s + p_s)\theta_s, \qquad \mathcal{A} = \sum_s (\rho_s + p_s)\sigma_s, \tag{11.22}$$

where all densities are expressed in present critical-density units and k is measured in H_0. The implemented momentum and shear equations are

$$\Phi' + \mathcal{H}\Psi = \frac{3}{2}\frac{a^2}{k^2}\mathcal{Q}, \tag{11.23}$$

$$\Phi - \Psi = \frac{9}{2}\frac{a^2}{k^2}\mathcal{A}. \tag{11.24}$$

The energy equation is not imposed after every step. It is retained as an independent residual,

$$\mathcal{C}_E := k^2\Phi + 3\mathcal{H}(\Phi' + \mathcal{H}\Psi) + \frac{3}{2}a^2\delta\rho. \tag{11.25}$$

The largest relative accepted-state residual in the main run is

$$\boxed{\epsilon_E = 7.5049074572 \times 10^{-4}.} \tag{11.26}$$

The dedicated refinement runs reduce the largest accepted-state residual to 1.2091040×10^{-4} on their common sample.

11.5 Adiabatic production-surface data

The perturbations are matched at $a = a_\star$ using one adiabatic mode. Let

$$\frac{\delta_s}{1+w_s} = C_\star \tag{11.27}$$

for every populated species, and choose $\Phi'_\star = 0$. With $\Phi_\star = \Psi_\star = 1$ as an intermediate normalization, the common velocity is

$$\theta_\star = \frac{k^2\mathcal{H}_\star\Psi_\star}{\frac{3}{2}a_\star^2\sum_s(\rho_s + p_s)}, \tag{11.28}$$

while the energy equation fixes

$$C_\star = -\frac{k^2\Phi_\star + 3\mathcal{H}_\star^2\Psi_\star}{\frac{3}{2}a_\star^2\sum_s(\rho_s + p_s)}. \tag{11.29}$$

Thus

$$\delta_b = C_\star, \qquad \delta_\gamma = \delta_\nu = \frac{4}{3}C_\star, \qquad \delta_D = (1+w_D)C_\star. \tag{11.30}$$

The source functions are finally divided by the comoving-curvature normalization

$$\mathcal{R}_\star = \Phi_\star - \frac{1}{3}C_\star. \tag{11.31}$$

This is a declared matching condition on the production hypersurface. The microscopic inhomogeneous crossing has not yet been solved, so the calculation does not claim that the full theory uniquely selects this adiabatic state.

11.6 Tight-coupling stage

From $a_\star$ to

$$a_{\rm sw} = 6.8 \times 10^{-4}, \tag{11.32}$$

photons and baryons are evolved as one tight-coupled velocity θ. Set

$$R_b(a) = \frac{3\rho_b}{4\rho_\gamma} = \frac{3\Omega_{b0}a}{4\Omega_{\gamma0}}. \tag{11.33}$$

The leading brightness moments are

$$F_{\gamma1} = \frac{4\theta}{3k}, \qquad F_{\gamma2} = \frac{32\theta}{45\dot\kappa}. \tag{11.34}$$

The common Euler equation is

$$\theta' = -\mathcal{H}\frac{R_b}{1+R_b}\theta + k^2\left[\frac{\delta_\gamma}{4(1+R_b)} + \Psi + \frac{R_b}{1+R_b}c_b^2\delta_b\right]. \tag{11.35}$$

The continuity equations are

$$\delta_b' = -\theta + 3\Phi', \qquad \delta_\gamma' = -\frac{4}{3}\theta + 4\Phi'. \tag{11.36}$$

At $a_{\rm sw}$, the photon quadrupole is matched into the full hierarchy, and the scalar polarization moments are initialized by

$$G_{\gamma0} = \frac{5}{4}F_{\gamma2}, \qquad G_{\gamma2} = \frac{1}{4}F_{\gamma2}. \tag{11.37}$$

This matching prescription is part of the controlled approximation and is tested through the multipole-refinement audit below.

11.7 Photon temperature and polarization hierarchy

Expand the photon temperature brightness and scalar polarization brightness as

$$F_\gamma(\mu) = \sum_{\ell\geq0}(-i)^\ell(2\ell+1)F_{\gamma\ell}P_\ell(\mu), \tag{11.38}$$

$$G_\gamma(\mu) = \sum_{\ell\geq0}(-i)^\ell(2\ell+1)G_{\gamma\ell}P_\ell(\mu). \tag{11.39}$$

The conventions are

$$F_{\gamma 0} = \delta_\gamma, \qquad F_{\gamma 1} = \frac{4\theta_\gamma}{3k}, \qquad F_{\gamma 2} = 2\sigma_\gamma. \tag{11.40}$$

Define the Thomson quadrupole source

$$\Pi_\gamma = F_{\gamma 2} + G_{\gamma 0} + G_{\gamma 2}. \tag{11.41}$$

The implemented temperature equations are

$$F'_{\gamma 0} = -kF_{\gamma 1} + 4\Phi', \tag{11.42}$$

$$F'_{\gamma 1} = \frac{k}{3}(F_{\gamma 0} - 2F_{\gamma 2}) + \frac{4k}{3}\Psi + \dot\kappa\left(\frac{4\theta_b}{3k} - F_{\gamma 1}\right), \tag{11.43}$$

$$F'_{\gamma 2} = \frac{2k}{5}\left(F_{\gamma 1} - \frac{3}{4}F_{\gamma 3}\right) - \dot\kappa\left(F_{\gamma 2} - \frac{1}{10}\Pi_\gamma\right), \tag{11.44}$$

$$F'_{\gamma \ell} = \frac{k}{2\ell+1}\left[\ell F_{\gamma,\ell-1} - (\ell+1)F_{\gamma,\ell+1}\right] - \dot\kappa F_{\gamma\ell}, \qquad \ell \geq 3. \tag{11.45}$$

The polarization equations are

$$G'_{\gamma 0} = -kG_{\gamma 1} - \dot\kappa\left(G_{\gamma 0} - \frac{1}{2}\Pi_\gamma\right), \tag{11.46}$$

$$G'_{\gamma 1} = \frac{k}{3}(G_{\gamma 0} - 2G_{\gamma 2}) - \dot\kappa G_{\gamma 1}, \tag{11.47}$$

$$G'_{\gamma 2} = \frac{2k}{5}\left(G_{\gamma 1} - \frac{3}{4}G_{\gamma 3}\right) - \dot\kappa\left(G_{\gamma 2} - \frac{1}{10}\Pi_\gamma\right), \tag{11.48}$$

with the collisionally damped free-streaming recurrence for $\ell \geq 3$.

The main solve uses

$$L_\gamma = L_E = 7. \tag{11.49}$$

At the angular boundary, a spherical-Bessel recurrence estimates the next multipole using the free-streaming phase accumulated from the tight-coupling switch.

11.8 Baryons and massless neutrinos

After the switch, baryons obey

$$\delta'_b = -\theta_b + 3\Phi', \tag{11.50}$$

$$\theta'_b = -\mathcal{H}\theta_b + c_b^2 k^2 \delta_b + \frac{4\rho_\gamma}{3\rho_b}\dot\kappa(\theta_\gamma - \theta_b) + k^2\Psi. \tag{11.51}$$

The controlled thermal model keeps baryons at the photon temperature through recombination and cools them adiabatically afterward. The baryon pressure term is numerically negligible for the primary acoustic dynamics but is retained.

Massless neutrinos use the collisionless brightness hierarchy

$$F'_{\nu 0} = -kF_{\nu 1} + 4\Phi', \tag{11.52}$$

$$F'_{\nu 1} = \frac{k}{3}(F_{\nu 0} - 2F_{\nu 2}) + \frac{4k}{3}\Psi, \tag{11.53}$$

$$F'_{\nu \ell} = \frac{k}{2\ell+1}\left[\ell F_{\nu,\ell-1} - (\ell+1)F_{\nu,\ell+1}\right], \qquad \ell \geq 2, \tag{11.54}$$

with

$$L_\nu = 6. \tag{11.55}$$

The terminal recurrence uses the total free-streaming phase accumulated from $a_\star$, not merely the post-switch phase. This distinction is required for stable early-time neutrino closure.

11.9 Massive refinement hierarchy

For every retained (h, q) node, write

$$\mathcal{F}_h(k, q, \mu, \eta) = \sum_{\ell \geq 0} (-i)^\ell (2\ell+1) F_{h\ell} P_\ell(\mu), \tag{11.56}$$

with

$$D_h(q) = \frac{d \log \bar{f}_h}{d \log q}, \qquad v_h(q, a) = \frac{q}{\epsilon_h(q, a)}. \tag{11.57}$$

The hierarchy is

$$F'_{h0} = -k v_h F_{h1} - \Phi' D_h, \tag{11.58}$$

$$F'_{h1} = \frac{k v_h}{3}(F_{h0} - 2F_{h2}) - \frac{k}{3 v_h}\Psi D_h, \tag{11.59}$$

$$F'_{h\ell} = \frac{k v_h}{2\ell+1}\left[\ell F_{h,\ell-1} - (\ell+1)F_{h,\ell+1}\right], \qquad \ell \geq 2. \tag{11.60}$$

The sign in (11.58) is fixed so that the massless limit $D_h = -4$ reproduces the photon gravitational source $+4\Phi'$.

The massive density, pressure, momentum, and shear moments are reconstructed from the same hierarchy. For example,

$$\delta\rho_D = \frac{1}{2\pi^2 a^4} \sum_h d_h \int q^2 \epsilon_h \bar{f}_h F_{h0}\, dq, \tag{11.61}$$

$$\mathcal{Q}_D = \frac{k}{2\pi^2 a^4} \sum_h d_h \int q^3 \bar{f}_h F_{h1}\, dq, \tag{11.62}$$

$$\mathcal{A}_D = \frac{1}{3\pi^2 a^4} \sum_h d_h \int \frac{q^4}{\epsilon_h} \bar{f}_h F_{h2}\, dq. \tag{11.63}$$

The main CMB solve uses $L_D = 2$. The refinement audit increases the photon, polarization, neutrino, and dark angular cutoffs simultaneously and finds a source discrepancy below one percent.

11.10 Numerical phase space and integration

The main deterministic state contains

component	resolution
wavenumbers	101
retained refinement levels	7
momentum nodes per level	2
total massive phase-space nodes	14
photon temperature multipoles	$0 \le \ell \le 7$
photon polarization multipoles	$0 \le \ell \le 7$
massless-neutrino multipoles	$0 \le \ell \le 6$
massive refinement multipoles	$0 \le \ell \le 2$

The equations are integrated in $N = \log a$ by fourth-order Runge–Kutta. The tight-coupling stage uses $\Delta N \simeq 3.5 \times 10^{-3}$ and the full collision stage uses $\Delta N \simeq 8.0 \times 10^{-4}$. Source records extend to

$$a_{\mathrm{end}} = 1.55 \times 10^{-3}, \tag{11.64}$$

which lies beyond the visibility peak.

The solved source functions contain visible acoustic oscillations over the complete k interval.

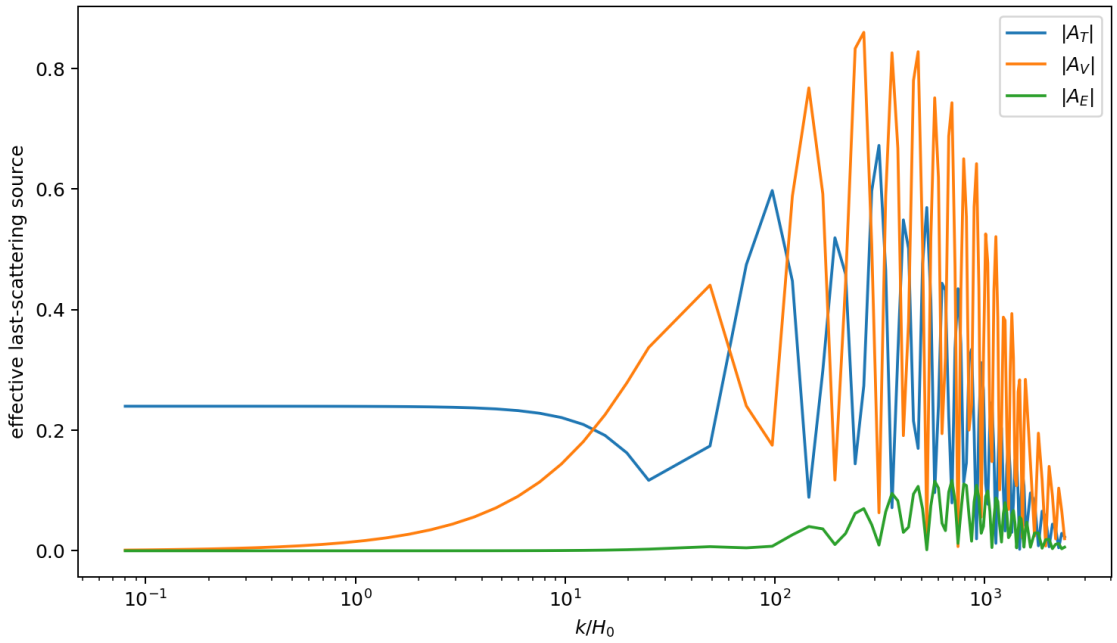


Figure 11.2. Effective last-scattering temperature, velocity, and polarization source amplitudes. The high-k oscillations are sampled on the 101-mode Einstein–Boltzmann grid and spline-interpolated only at the final projection stage.

11.11 Finite-width last-scattering projection

Define

$$\Theta_0 = \frac{1}{4} F_{\gamma 0}. \tag{11.65}$$

The source amplitudes in dN measure are

$$S_T = g_N \left(\Theta_0 + \Psi + \frac{1}{16}\Pi_\gamma \right) + e^{-\tau}(\Psi_N - \Phi_N), \tag{11.66}$$

$$S_V = g_N \frac{\theta_b}{k}, \tag{11.67}$$

$$S_E = \frac{3}{16} g_N \Pi_\gamma. \tag{11.68}$$

Their normalized integrated amplitudes are

$$A_X(k) = \frac{1}{\mathcal{R}_\star} \int S_X(k, N)\, dN, \qquad X \in \{T, V, E\}. \tag{11.69}$$

Let the visibility-weighted distance and width be

$$\chi_\star = \frac{\int g_N \chi\, dN}{\int g_N\, dN}, \qquad \sigma_\chi^2 = \frac{\int g_N (\chi - \chi_\star)^2 dN}{\int g_N dN}. \tag{11.70}$$

The numerical values are

$$\boxed{H_0 \chi_\star = 3.1438585561, \qquad H_0 \sigma_\chi = 2.3136123623 \times 10^{-3}.} \tag{11.71}$$

The finite-width projection used here is

$$\Delta_\ell^T(k) = e^{-k^2\sigma_\chi^2/2} \left[A_T(k) j_\ell(k\chi_\star) + A_V(k) j_\ell'(k\chi_\star) \right], \tag{11.72}$$

$$\Delta_\ell^E(k) = e^{-k^2\sigma_\chi^2/2} \sqrt{\frac{(\ell+2)!}{(\ell-2)!}} A_E(k) \frac{j_\ell(k\chi_\star)}{(k\chi_\star)^2}. \tag{11.73}$$

This is a controlled narrow-source approximation. A precision implementation would retain the complete time dependence inside the line-of-sight integral rather than compressing the source to $(A_T, A_V, A_E, \chi_\star, \sigma_\chi)$.

For

$$\mathcal{P}_\mathcal{R}(k) = A_s \left(\frac{k}{k_\star} \right)^{n_s - 1}, \tag{11.74}$$

with $A_s = 2.1 \times 10^{-9}$, $n_s = 0.965$, and $k_\star = 0.05\,\mathrm{Mpc}^{-1}$, the angular spectra are

$$C_\ell^{XY} = 4\pi \int d\log k\, \mathcal{P}_\mathcal{R}(k) \Delta_\ell^X(k) \Delta_\ell^Y(k), \qquad X, Y \in \{T, E\}. \tag{11.75}$$

11.12 Acoustic and polarization results

The visibility-peak sound horizon is

$$H_0 r_s = 0.02913803155, \tag{11.76}$$

and the geometric acoustic ruler is

$$\boxed{\ell_A = \pi \frac{\chi_{\rm vis}}{r_s} = 339.055741.} \tag{11.77}$$

The first three resolved TT acoustic maxima above $\ell = 80$ occur at

$$\boxed{\ell_1 = 252, \qquad \ell_2 = 608, \qquad \ell_3 = 904.} \tag{11.78}$$

Their spacing is consistent with the geometric ruler to a relative discrepancy of 4.9975×10^{-2} in the declared projection.

The first peak has

$$\boxed{D_{\ell_1}^{TT} := \frac{\ell_1(\ell_1+1)}{2\pi} C_{\ell_1}^{TT} T_{\gamma 0}^2 = 6163.119890\,\mu{\rm K}^2.} \tag{11.79}$$

Selected template values are

ℓ	D_ℓ^{TT} [μK^2]	D_ℓ^{EE} [μK^2]	D_ℓ^{TE} [μK^2]
100	1220.174	0.289	-1.418
200	4733.471	0.428	-16.260
252	6163.120	0.459	17.919
500	2395.207	8.980	-45.853
608	3467.906	4.787	27.666
904	3099.722	9.410	-46.574
1000	1904.695	11.920	58.657
1500	343.904	15.671	-25.350
2000	138.183	1.915	2.085
2500	10.197	1.762	-0.480

The temperature–polarization cross spectrum satisfies the Cauchy bound over every stored multipole,

$$\max_\ell \frac{|C_\ell^{TE}|}{\sqrt{C_\ell^{TT} C_\ell^{EE}}} = 0.9948225265 < 1. \tag{11.80}$$

The finite source width suppresses high multipoles. At $\ell = 3000$, the declared Gaussian thickness factor is

$$\exp\left[-\frac{1}{2}\left(\frac{\ell \sigma_\chi}{\chi_\star}\right)^2\right]_{\ell=3000} = 0.08741675187. \tag{11.81}$$

11.13 Convergence and verification

The main deterministic audit contains

$$\boxed{\text{20 PASS,} \qquad \text{0 FAIL.}} \tag{11.82}$$

A separate convergence audit adds

$$\boxed{\text{4 PASS,} \qquad \text{0 FAIL.}} \tag{11.83}$$

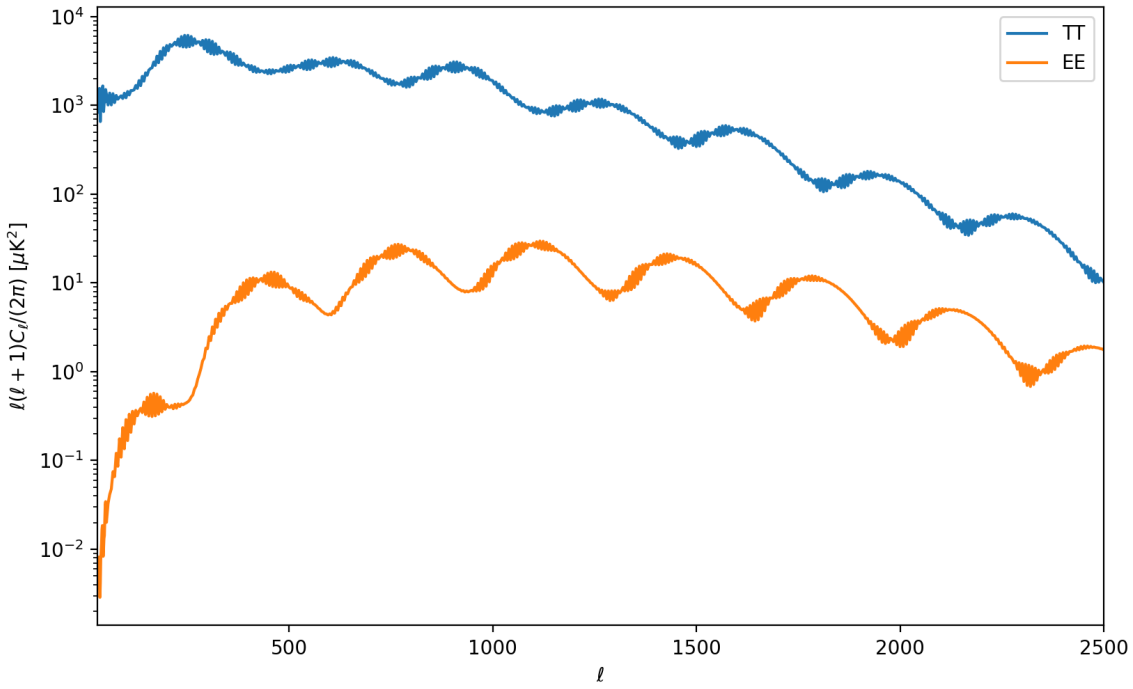


Figure 11.3. Controlled recombination-era TT and EE templates. The plotted range begins at $\ell = 30$ because reionization and late integrated sources are not included.

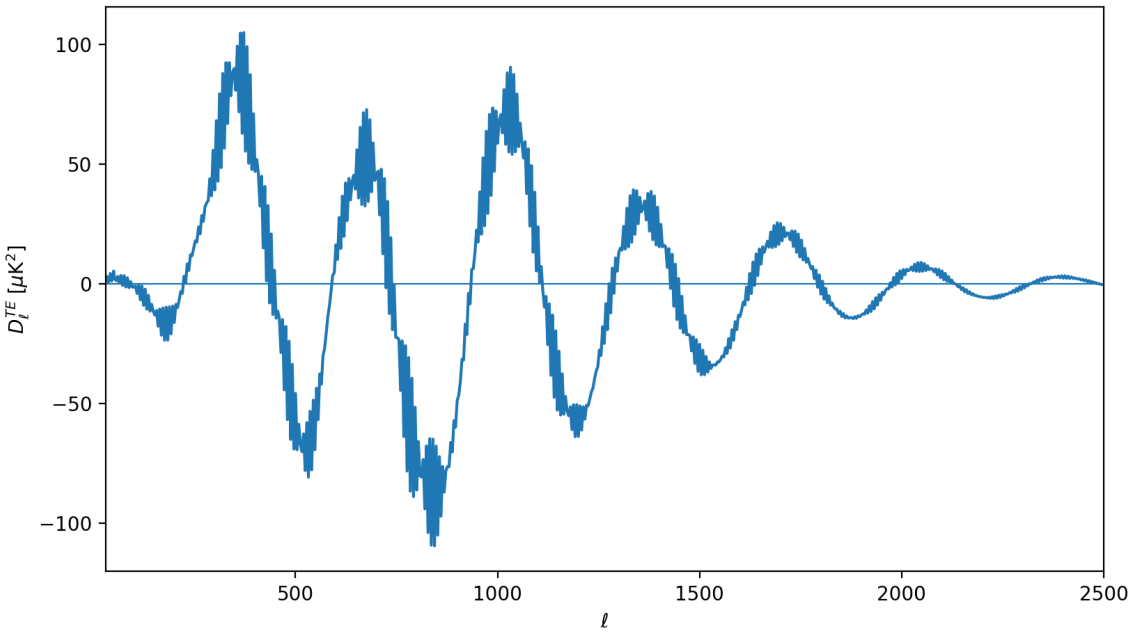


Figure 11.4. The controlled TE template alternates sign across the acoustic series, reflecting the relative phase of photon compression, velocity, and polarization.

The combined calculation therefore records

$$\boxed{\text{24 PASS,} \qquad \text{0 FAIL.}} \tag{11.84}$$

The principal refinement differences are

control	relative discrepancy
momentum quadrature	1.81645×10^{-4}
photon–neutrino–dark multipole truncation	6.48202×10^{-3}
time step	2.19752×10^{-6}
accepted-state Einstein constraint	1.20910×10^{-4}

The free-streaming spherical-Bessel identity is reproduced with maximum residual

$$\epsilon_{\mathrm{fs}} = 1.35046 \times 10^{-6} \tag{11.85}$$

in the finite-difference derivative audit. Median held-out source-interpolation errors are

$$\epsilon_{A_T} = 1.38235 \times 10^{-2}, \qquad \epsilon_{A_V} = 1.45514 \times 10^{-2}, \qquad \epsilon_{A_E} = 1.42026 \times 10^{-2}. \tag{11.86}$$

These source errors, together with the narrow-source projection, dominate over the time-step and momentum-quadrature errors in the current spectra.

11.14 Closure ledger

statement	status	boundary
photon temperature hierarchy with Thomson scattering	closed controlled	scalar linear perturbations
scalar polarization hierarchy and nonzero E source	closed controlled	no B modes
baryon drag and acoustic loading	closed controlled	simple baryon thermal history
massless-neutrino hierarchy and anisotropic stress	closed controlled	fixed $\Omega_{\nu 0}$
massive refinement hierarchy through recombination	closed controlled	seven levels, fourteen phase nodes
Einstein momentum and shear coupling	closed controlled	Newtonian gauge
energy-constraint audit	numerically controlled	not projected after each step
Saha visibility and finite source width	closed conditional	no multilevel recombination
recombination-era TT, TE, EE acoustic templates	numerically closed	narrow-source projection
low-ℓ TT, TE, EE prediction	open	reionization and late integrated sources
precision recombination and helium kinetics	open	external atomic network
full line-of-sight source integration	open	production-grade transfer solver
lensing remapping and nonlinear late-time sources	open	nonlinear structure
joint likelihood and parameter inference	open	observational data and nuisance model
microscopic perturbation matching at production	open	nonlinear inhomogeneous crossing

11.15 Synthesis and scope

The scale-covariant distribution now enters one coupled chain,

$$\boxed{\begin{aligned}&\bar{f}_h(q) \longrightarrow \{F_{h\ell}, F_{\nu\ell}, F_{\gamma\ell}, G_{\gamma\ell}, \delta_b, \theta_b\}\\ &\longrightarrow \{\Phi, \Psi, g_N\} \longrightarrow \{A_T, A_V, A_E\} \longrightarrow \{C_\ell^{TT}, C_\ell^{TE}, C_\ell^{EE}\}.\end{aligned}} \tag{11.87}$$

No dark-sector parameter is changed between production, background expansion, late-time collisionless structure, and the controlled primary-CMB calculation.

This closes the principal species-coupling gap at the level of a deterministic linear template. A precision line-of-sight and thermal-history treatment still requires:

$$\boxed{\begin{aligned}&\text{replace Saha recombination by multilevel H/He kinetics,}\\ &\text{derive the production-surface perturbation transfer,}\\ &\text{add reionization, late ISW, and lensing remapping,}\\ &\text{compute primary and late-time observables in one solver,}\\ &\text{then perform a first joint likelihood without dark-sector retuning.}\end{aligned}} \tag{11.88}$$

CHAPTER 12

Effective Thermal History, Production-Surface Matching, Reionization, and Late-Time Line-of-Sight Closure

12.1 Scope and objective

The preceding production-to-primary-CMB calculation jointly evolved photons, scalar polarization, baryons, massless neutrinos, metric perturbations, and the massive refinement hierarchy through a controlled recombination surface. Its stated thermal and projection limitations were Saha hydrogen recombination, no reionization, no late integrated Sachs–Wolfe source, and a narrow finite-width last-scattering projection. This chapter extends that framework by adding

$$\boxed{\begin{array}{c}\text{effective Peebles hydrogen recombination} + \text{ controlled helium transitions} \\ + \text{ optical-depth-calibrated reionization} + \text{ single-clock production matching} \\ + \text{ finite-width recombination and reionization projection} + \text{ time-resolved late ISW transport.}\end{array}} \tag{12.1}$$

The early acoustic source amplitudes are inherited from the coupled Einstein–Boltzmann solve without dark-sector retuning. The new calculation is a controlled thermal-history and late-source closure. It is not a precision recombination code, an astrophysical reionization prediction, or a likelihood analysis.

A sign correction is also made. In the Newtonian-gauge convention used here, the integrated temperature source is proportional to

$$\Phi' + \Psi', \tag{12.2}$$

not to their difference. The corrected sum is used throughout this chapter.

12.2 Unchanged cosmological and refinement input

Use

$$E(a)^2 = \Omega_{r0} a^{-4} + \Omega_{m0} a^{-3} + \Omega_{V0}, \qquad \Omega_{m0} = \Omega_{b0} + \Omega_{D0}, \tag{12.3}$$

with

$$\Omega_{b0} = 0.049, \qquad \Omega_{D0} = 0.261, \qquad \Omega_{r0} = 9 \times 10^{-5}, \qquad \Omega_{V0} = 1 - \Omega_{m0} - \Omega_{r0}. \tag{12.4}$$

The conditional physical conversion uses

$$H_0 = 70\ \mathrm{km\,s^{-1}\,Mpc^{-1}}, \qquad T_{\gamma 0} = 2.7255\ mK, \qquad Y_p = 0.24. \tag{12.5}$$

The primordial spectrum remains

$$\mathcal{P}_{\mathcal{R}}(k) = A_s \left(\frac{k}{k_\star}\right)^{n_s-1}, \qquad A_s = 2.1 \times 10^{-9}, \qquad n_s = 0.965, \qquad k_\star = 0.05\,\mathrm{Mpc}^{-1}. \tag{12.6}$$

The recombination-era acoustic source functions

$$A_T(k), \qquad A_V(k), \qquad A_E(k) \tag{12.7}$$

are read directly from the preceding coupled solve. No parameter of the scale-covariant dark sector is refitted in this calculation.

12.3 Effective Peebles hydrogen recombination

Let

$$N = \log a, \qquad T_\gamma = \frac{T_{\gamma 0}}{a}, \qquad n_H(a) = n_{H0} a^{-3}. \tag{12.8}$$

The case-B coefficient is represented by the controlled fit

$$\alpha_B(T) = 4.309 \times 10^{-19} \frac{t^{-0.6166}}{1 + 0.6703 t^{0.5300}}\ \mathrm{m^3\,s^{-1}}, \qquad t = \frac{T}{10^4\ \mathrm{K}}. \tag{12.9}$$

Detailed balance gives

$$\beta_B(T) = \alpha_B(T) \left(\frac{m_e k_B T}{2\pi\hbar^2}\right)^{3/2} \exp\left(-\frac{E_2}{k_B T}\right), \qquad E_2 = 3.4\ \mathrm{eV}. \tag{12.10}$$

Define

$$K_\alpha = \frac{\lambda_\alpha^3}{8\pi H}, \qquad C_P = \frac{1 + K_\alpha \Lambda_{2s} n_H (1 - x_H)}{1 + K_\alpha (\Lambda_{2s} + \beta_B) n_H (1 - x_H)}. \tag{12.11}$$

The effective three-level hydrogen equation is

$$\boxed{\frac{dx_H}{dN} = \frac{C_P}{H}\left[\beta_B(1 - x_H) \exp\left(-\frac{E_\alpha}{k_B T}\right) - n_H \alpha_B x_H^2\right],} \tag{12.12}$$

where

$$E_\alpha = 10.2\ \mathrm{eV}, \qquad \Lambda_{2s} = 8.22458\ \mathrm{s^{-1}}, \qquad \lambda_\alpha = 121.567\ \mathrm{nm}. \tag{12.13}$$

The calculation follows the Saha solution until

$$x_H \simeq 0.99, \qquad z \simeq 1577.09, \tag{12.14}$$

and then integrates (12.12). It gives

$$\boxed{x_{H,\mathrm{residual}} = 2.8029283896 \times 10^{-4}.} \tag{12.15}$$

The helium contribution is represented by controlled smooth transitions for HeIII-to-HeII and HeII-to-HeI. This preserves the correct electron-counting limits but is not a multilevel helium-kinetics calculation.

12.4 Reionization and optical-depth calibration

Let

$$y(z) = (1+z)^{3/2}. \tag{12.16}$$

The hydrogen reionization fraction is modeled by

$$R_H(z) = \frac{1}{2}\left[1 + \tanh\left(\frac{y(z_{\rm re}) - y(z)}{\Delta y_{\rm re}}\right)\right], \tag{12.17}$$

with

$$\Delta y_{\rm re} = \frac{3}{2}\sqrt{1+z_{\rm re}}\,\Delta z_{\rm re}, \qquad \Delta z_{\rm re} = 0.5. \tag{12.18}$$

A second smooth transition accounts for the late second helium ionization.

The optical-depth equation in logarithmic scale factor is

$$\frac{d\tau}{dN} = \frac{n_e \sigma_T c}{H}, \tag{12.19}$$

and the visibility per unit N is

$$g_N(N) = \exp[-\tau(N)]\frac{d\tau}{dN}. \tag{12.20}$$

The reionization midpoint is fixed by the declared input

$$\tau_{\rm re} = 0.054. \tag{12.21}$$

The calibrated result is

$$\boxed{z_{\rm re} = 7.4124651067,} \qquad \boxed{z_{g,\rm re} = 6.6559660691.} \tag{12.22}$$

The total visibility satisfies

$$\boxed{\int g_N\, dN = 0.9999979519.} \tag{12.23}$$

The recombination and reionization pieces carry

$$\int g_{N,\rm rec}\, dN = 0.9474304116, \qquad \int g_{N,\rm re}\, dN = 0.0525675402. \tag{12.24}$$

12.5 Visibility moments and finite-thickness projection

For a nonnegative source weight $w(N)$, define

$$\mathcal{N}_w = \int w\, dN, \qquad \chi_w = \frac{1}{\mathcal{N}_w}\int w\chi\, dN, \tag{12.25}$$

$$\sigma^2_{\chi,w} = \frac{1}{\mathcal{N}_w}\int w(\chi - \chi_w)^2\, dN. \tag{12.26}$$

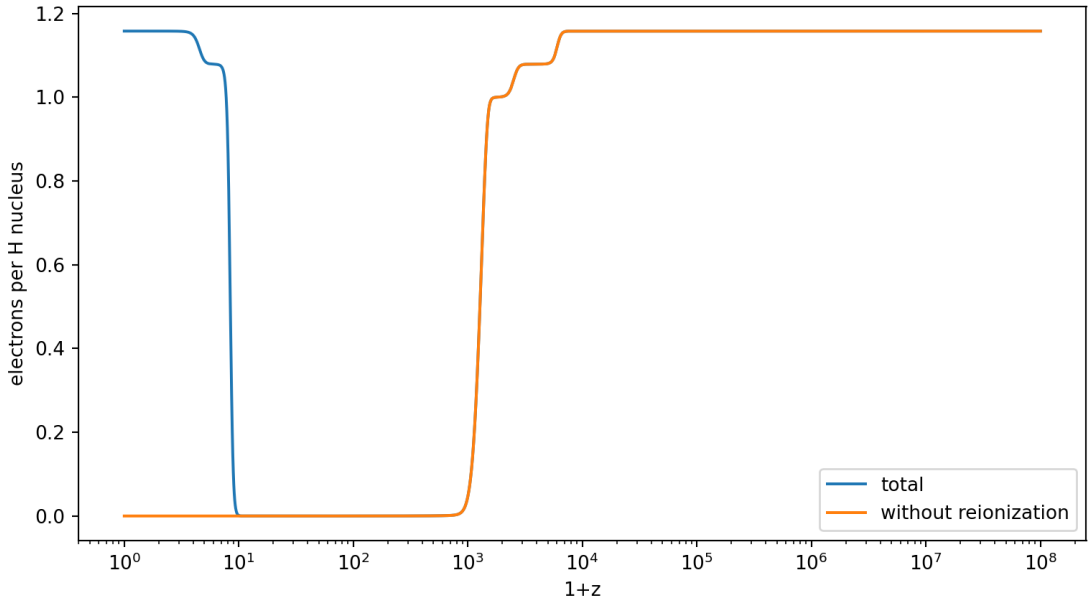


Figure 12.1. Effective hydrogen-plus-helium ionization history before and after the optical-depth-calibrated reionization transition.

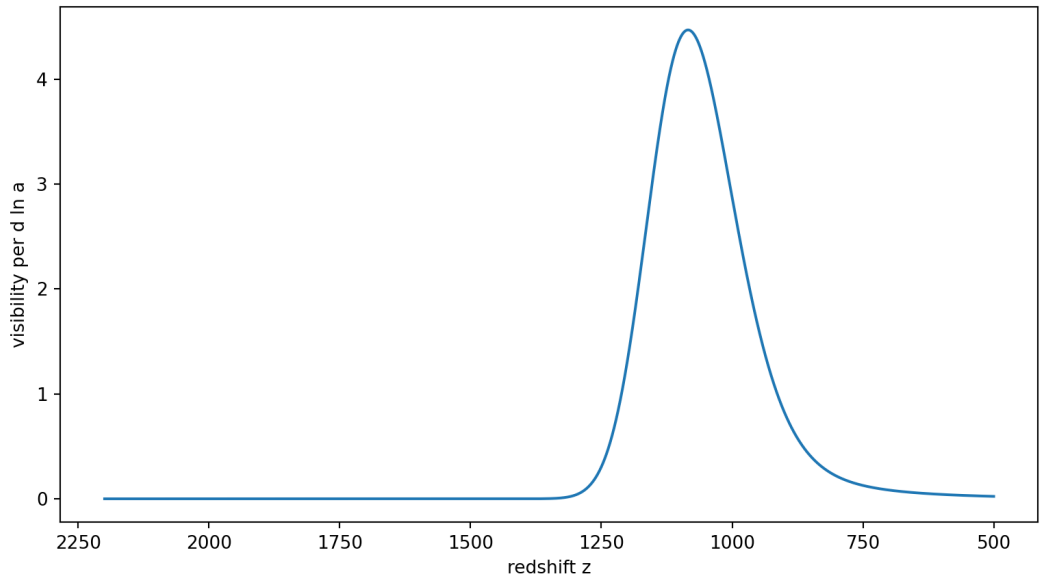


Figure 12.2. Effective recombination visibility. The peak shifts from the preceding Saha benchmark to the controlled Peebles value near $z = 1084$.

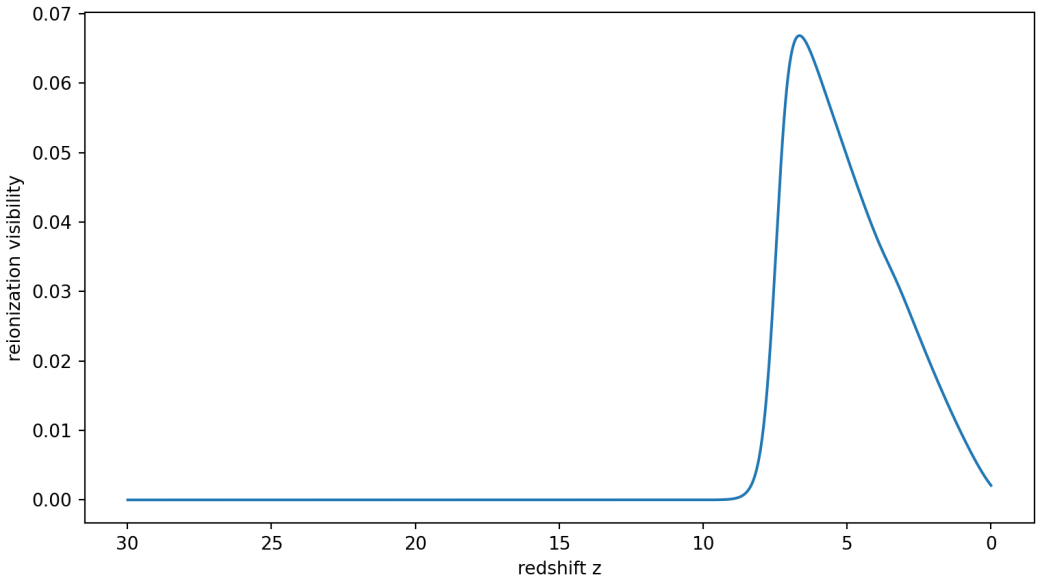


Figure 12.3. Reionization-only visibility after subtracting the recombination history.

The long residual-ionization tail makes the global recombination RMS width a poor Gaussian approximation to the acoustic core. The calculation therefore archives the full visibility and both widths, while the acoustic thickness factor uses the FWHM-equivalent central width

$$\sigma_{\chi,\mathrm{core}} = \frac{\Delta\chi_{\mathrm{FWHM}}}{2\sqrt{2\log 2}}. \tag{12.27}$$

The measured quantities are

$$\boxed{\chi_{\mathrm{rec}} = 3.1335298685\ H_0^{-1}, \qquad \sigma_{\chi,\mathrm{core}} = 3.7505054558 \times 10^{-3}\ H_0^{-1},} \tag{12.28}$$

$$\sigma_{\chi,\mathrm{RMS}} = 1.6282448527 \times 10^{-2}\ H_0^{-1}, \tag{12.29}$$

$$\boxed{\chi_{\mathrm{re}} = 1.5792968823\ H_0^{-1}, \qquad \sigma_{\chi,\mathrm{re}} = 0.4275528337\ H_0^{-1}.} \tag{12.30}$$

The recombination visibility peaks at

$$\boxed{z_{\mathrm{rec}} = 1084.47305577.} \tag{12.31}$$

12.6 Single-clock production-surface matching

Let the production hypersurface be defined by one local clock F,

$$F(t, \boldsymbol{x}) = F_\star. \tag{12.32}$$

At first perturbative order, its displacement is

$$\delta t_\star = -\frac{\delta F}{\dot{\bar{F}}}. \tag{12.33}$$

Suppose every produced refinement level depends on the same clock and no independent modulation field is present. Then

$$\delta\rho_h = -\dot{\bar{\rho}}_h\,\delta t_\star. \tag{12.34}$$

The curvature perturbation of each level is

$$\zeta_h = -\Phi - H\frac{\delta\rho_h}{\dot{\bar{\rho}}_h} = -\Phi + H\delta t_\star, \tag{12.35}$$

so that

$$\boxed{S_{hj} := 3(\zeta_h - \zeta_j) = 0.} \tag{12.36}$$

Thus a single-clock crossing supplies adiabatic production-surface matching at leading gradient order.

An independent modulation $A_\star(\boldsymbol{x})$ generates a calculable entropy direction. For

$$n_h \propto \nu_h^{3/2}\exp\left(-\frac{\pi m_h^2}{\nu_h}\right), \qquad \nu_h = A_\star m_h, \tag{12.37}$$

its response is

$$\boxed{\frac{\partial\log n_h}{\partial\log A_\star} = \frac{3}{2} + \frac{\pi m_h}{A_\star}.} \tag{12.38}$$

This distinguishes the adiabatic single-clock prediction from an independently modulated production event.

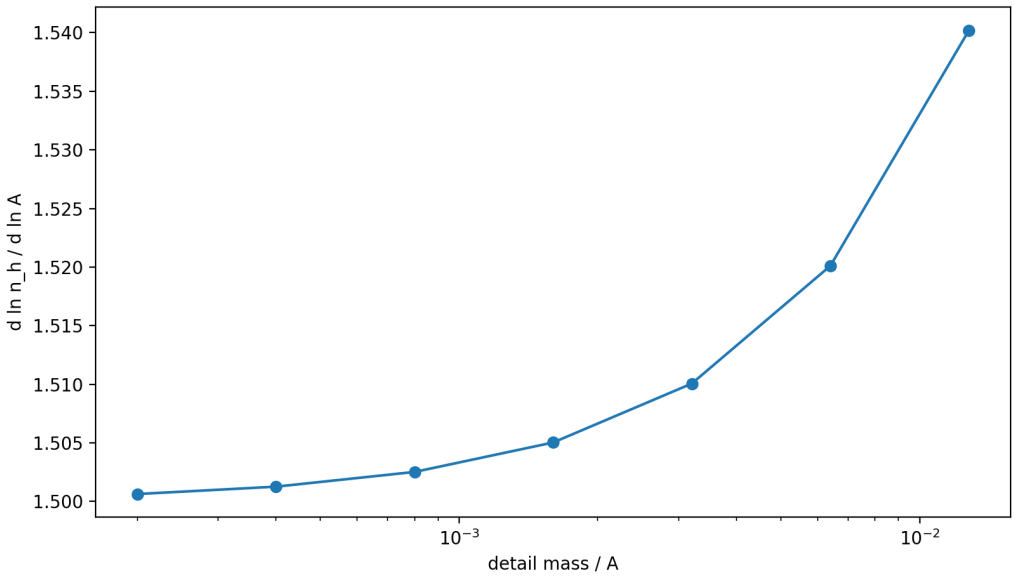


Figure 12.4. Level-dependent abundance response to an independent production-scale modulation. The single-clock relative entropy remains zero; the plotted response applies only to an additional modulation direction.

12.7 Corrected line-of-sight source

Use Newtonian gauge

$$ds^2 = a^2(\eta)\left[-(1+2\Psi)d\eta^2 + (1-2\Phi)d\boldsymbol{x}^2\right]. \tag{12.39}$$

For the collisionless temperature brightness after the standard gravitational redshift rearrangement,

$$(\Theta+\Psi)' + ik\mu(\Theta+\Psi) = \Phi' + \Psi' + \mathcal{C}_\gamma. \tag{12.40}$$

Therefore the late integrated source is

$$\boxed{S_{\rm ISW} = e^{-\tau}(\Phi'+\Psi').} \tag{12.41}$$

The earlier difference source is superseded in this convention.

Let

$$\mathcal{T}_\ell(k) = \mathcal{T}_{\ell,\rm rec}(k) + \mathcal{T}_{\ell,\rm re}(k) + \mathcal{T}_{\ell,\rm ISW}(k). \tag{12.42}$$

The finite-width recombination term is

$$\mathcal{T}_{\ell,\rm rec} = e^{-\tau_{\rm re}} \exp\left(-\frac{k^2\sigma^2_{\chi,\rm core}}{2}\right)\left[A_T(k) j_\ell(k\chi_{\rm rec}) + A_V(k) j'_\ell(k\chi_{\rm rec})\right]. \tag{12.43}$$

Free streaming from recombination to the effective reionization shell reconstructs

$$\Theta^{\rm re}_0(k) = A_T j_0(k\Delta\chi) + A_V j'_0(k\Delta\chi), \tag{12.44}$$

$$\Theta^{\rm re}_2(k) = A_T j_2(k\Delta\chi) + A_V j'_2(k\Delta\chi), \qquad \Pi_{\rm re} = 4\Theta^{\rm re}_2, \tag{12.45}$$

where

$$\Delta\chi = \chi_{\rm rec} - \chi_{\rm re}. \tag{12.46}$$

The reionization temperature and polarization sources are then projected with the measured reionization width.

The late potential is evolved with

$$D_{NN} + \left(2 + \frac{d\log H}{dN}\right) D_N - \frac{3}{2}\Omega_m(a) D = 0, \tag{12.47}$$

$$g_\Phi(a) = \frac{D(a)}{a}, \qquad \Phi = \Psi = \frac{3}{5}\mathcal{R}\, T(k) g_\Phi(a), \tag{12.48}$$

which gives the time-resolved integral

$$\mathcal{T}_{\ell,\rm ISW}(k) = \int dN\, 2e^{-\tau}\frac{d\Phi}{dN} j_\ell[k\chi(N)]. \tag{12.49}$$

The E-mode transfer is constructed from the recombination and reionization quadrupole sources with the scalar spin-2 projection factor

$$\sqrt{\frac{(\ell+2)!}{(\ell-2)!}}\frac{j_\ell(x)}{x^2}. \tag{12.50}$$

Finally,

$$\boxed{C_\ell^{XY} = 4\pi \int d\log k\, \mathcal{P}_\mathcal{R}(k) \mathcal{T}_\ell^X(k) \mathcal{T}_\ell^Y(k), \qquad X, Y \in \{T, E\}.} \tag{12.51}$$

12.8 Acoustic ruler and controlled spectra

The effective sound horizon is

$$r_s = \int^{N_{\rm rec}} \frac{dN}{\mathcal{H}\sqrt{3(1+R_b)}}, \qquad R_b = \frac{3\rho_b}{4\rho_\gamma}. \tag{12.52}$$

The calculation gives

$$\boxed{H_0 r_s = 0.03280490608,} \qquad \boxed{\ell_A = \frac{\pi \chi_{\rm rec}}{r_s} = 300.0854321.} \tag{12.53}$$

The first three resolved temperature maxima above $\ell = 150$ occur at

$$\boxed{\ell_1 = 262, \qquad \ell_2 = 582, \qquad \ell_3 = 894.} \tag{12.54}$$

The first peak amplitude is

$$\boxed{D_{262}^{TT} = 4063.713734\ \mu\mathrm{K}^2.} \tag{12.55}$$

ℓ	D_ℓ^{TT} [μK^2]	D_ℓ^{EE} [μK^2]	D_ℓ^{TE} [μK^2]
2	316.421	0.03564	1.6379
5	317.116	0.00992	0.9739
10	396.325	0.00006	0.1210
50	977.218	0.02039	1.2588
100	1279.639	0.17986	-1.0717
202	3277.910	0.33168	-11.7029
262	4063.714	0.55156	21.8382
498	1387.237	6.49005	-37.9429
582	1775.329	2.47052	-11.0646
894	946.625	3.64282	-24.6961
1002	533.522	4.39903	25.7940
1498	32.519	1.66852	-3.6807
2002	2.447	0.03852	0.1136

Table 12.1. Selected controlled thermal-history, reionization, and late-source spectra. These are templates rather than likelihood-grade spectra.

The reionization calculation produces a low-ℓ polarization enhancement. At $\ell = 5$,

$$\frac{D_{5,\rm full}^{EE}}{D_{5,\rm early}^{EE}} = 2927.1921, \tag{12.56}$$

because the early-only scalar E source is extremely small at that multipole. The absolute value remains controlled,

$$D_5^{EE} = 9.921 \times 10^{-3}\ \mu\mathrm{K}^2. \tag{12.57}$$

At $\ell = 10$, the late sources change the temperature template by

$$\left| \frac{D^{TT}_{10,\mathrm{full}}}{D^{TT}_{10,\mathrm{early}}} - 1 \right| = 0.5922236. \tag{12.58}$$

The TE Cauchy ratio obeys

$$\boxed{\max_{\ell} \frac{|C^{TE}_{\ell}|}{\sqrt{C^{TT}_{\ell} C^{EE}_{\ell}}} = 0.8112607541 < 1.} \tag{12.59}$$

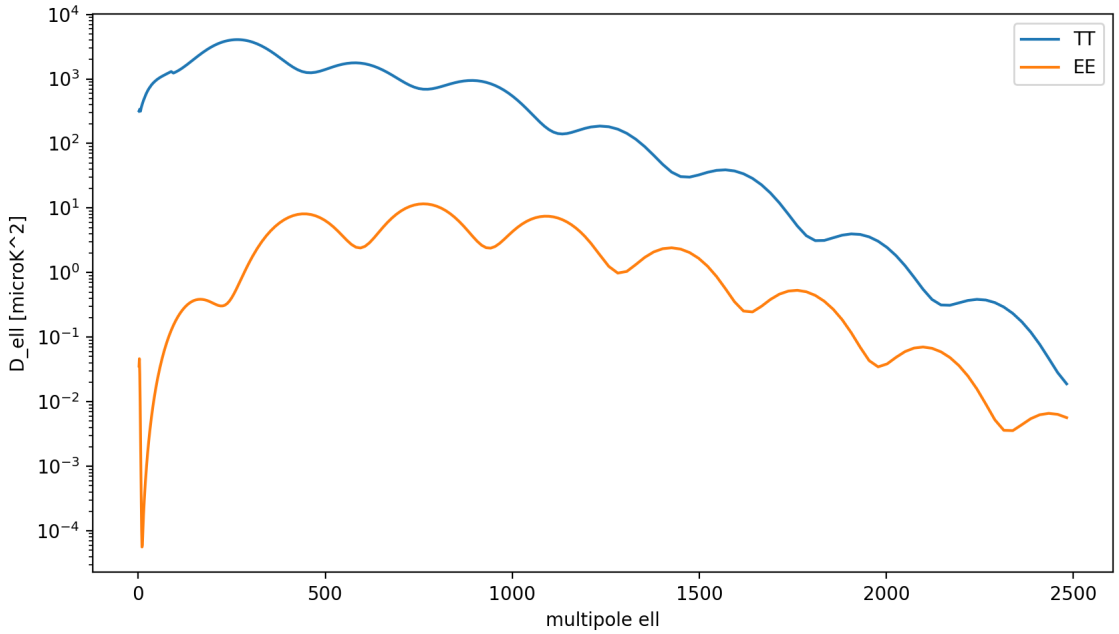


Figure 12.5. Controlled temperature and scalar E-mode templates after effective recombination and reionization.

12.9 Projection convergence and deterministic audit

The main projection uses 3350 global wavenumber nodes after combining logarithmic low-k sampling with a linear oscillatory grid. A 4500-node reference calculation was performed at selected multipoles.

control	measured maximum	acceptance bound
TT selected-grid discrepancy	1.20461×10^{-2}	6×10^{-2}
EE selected-grid discrepancy	4.42685×10^{-2}	6×10^{-2}
TE correlation discrepancy	7.86181×10^{-3}	6×10^{-2}
acoustic TT discrepancy	1.45217×10^{-5}	5×10^{-2}

Table 12.2. Main-to-fine projection convergence.

The main deterministic suite records

$$\boxed{\text{20 PASS,} \qquad \text{0 FAIL,}} \tag{12.60}$$

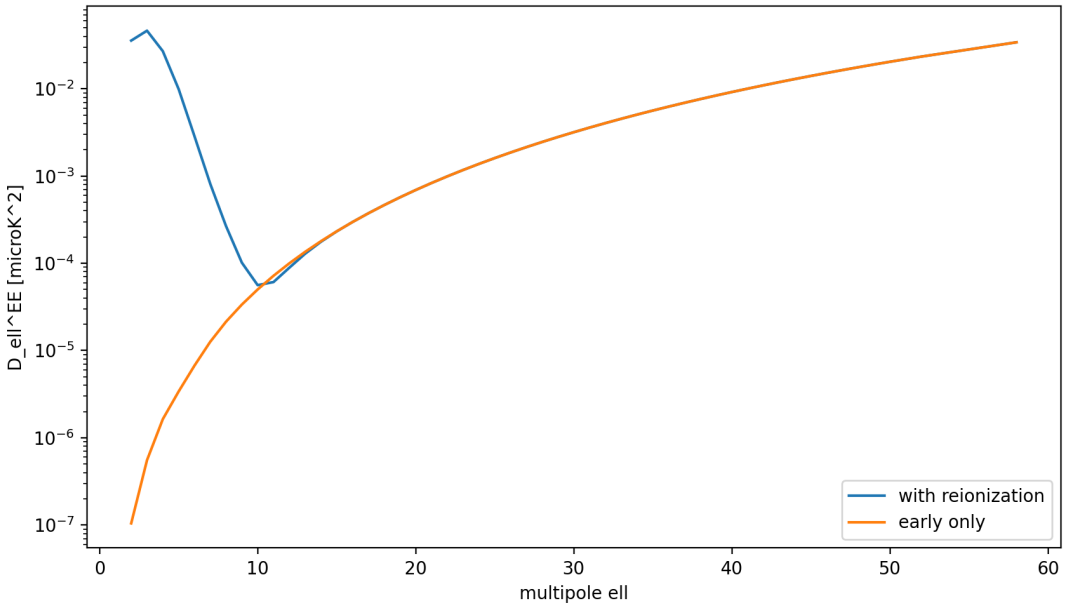


Figure 12.6. Reionization creates a low-multipole E-mode feature absent from the early-only source.

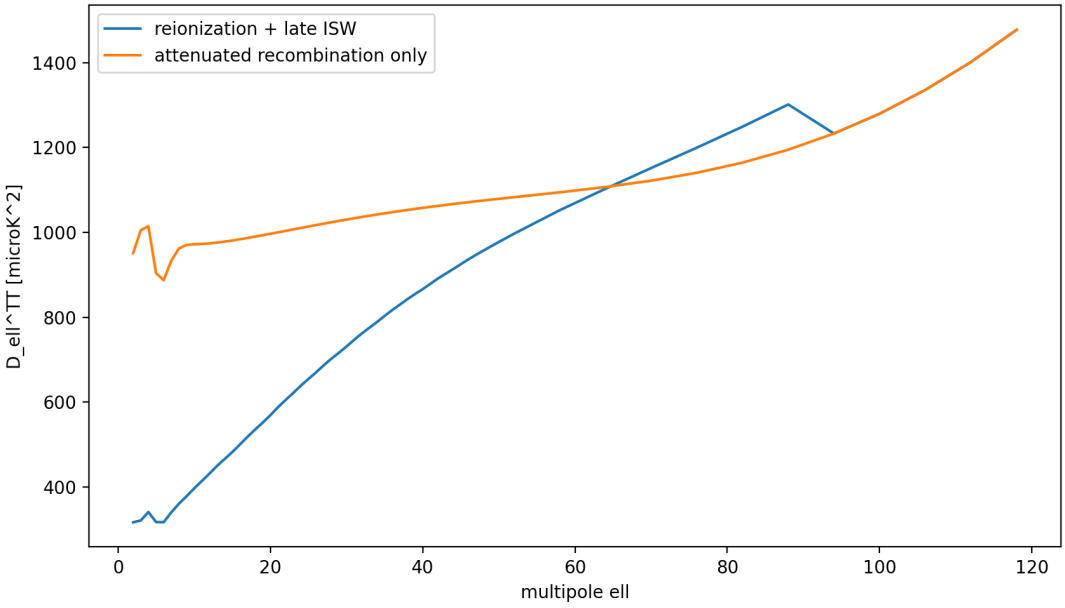


Figure 12.7. Late ISW and reionization modify the low-ℓ temperature source while leaving the acoustic regime effectively unchanged in this controlled projection.

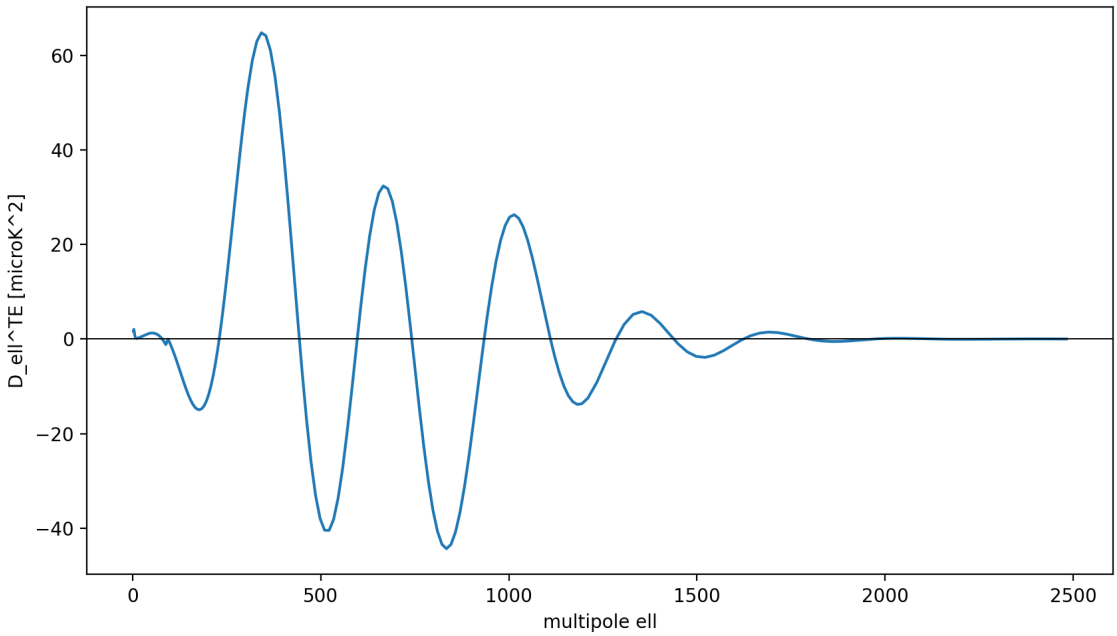


Figure 12.8. Controlled TE template with alternating correlation sign across the acoustic structure.

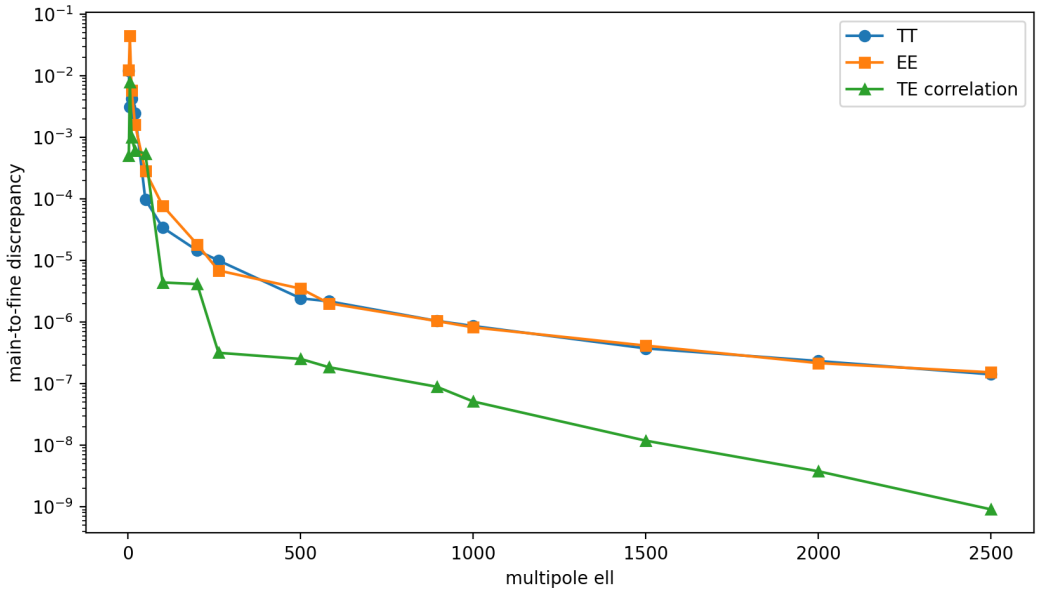


Figure 12.9. Selected-multipole convergence under global wavenumber-grid refinement.

and the independent projection-refinement suite records

$$\boxed{4 \text{ PASS}, \qquad 0 \text{ FAIL.}} \tag{12.61}$$

The combined result is

$$\boxed{24 \text{ PASS}, \qquad 0 \text{ FAIL.}} \tag{12.62}$$

12.10 Closure ledger

statement	status	boundary
effective Peebles hydrogen history	closed controlled	three-level approximation
controlled helium electron counting	closed conditional	smooth transition model
reionization optical depth	calibrated	τ_{re} remains an input
recombination and reionization visibility	closed controlled	declared background
single-clock production matching	proved at leading order	no independent modulation
independent modulation response	derived	linear response
late ISW source sign	corrected	stated Newtonian-gauge convention
late time-resolved ISW projection	closed controlled	linear potential growth
finite-width recombination projection	closed controlled	central-shell approximation
low-ℓ reionization polarization	demonstrated	effective source reconstruction
precision multilevel H/He recombination	open	atomic kinetics
reionization redshift from astrophysics	open	source population
full recombination-time LOS integral	open	shell approximation retained
CMB lensing remapping	open	nonlinear structure
joint likelihood and posterior	open	data interface

12.11 Synthesis and scope

This chapter closes three previously separate gaps:

$$\boxed{\begin{array}{c}\text{production-surface adiabatic matching,}\\ \text{effective recombination plus reionization,}\\ \text{corrected late line-of-sight temperature and polarization sources.}\end{array}} \tag{12.63}$$

It also separates a robust statement from a model input. A single-clock production event is adiabatic at leading gradient order, while an independent modulation field generates a calculable level-dependent entropy response. The reionization optical depth is calibrated rather than predicted.

CHAPTER 13

CMB Lensing Remapping, Lensing B Modes, and Cross-Observable Scale Separation

13.1 Scope and objective

The production-to-CMB calculation supplies controlled unlensed TT, TE, and EE templates, while the collisionless refinement calculation supplies a late-time scale-dependent matter transfer. The missing linear-observable link was the gravitational remapping between them. This chapter constructs a common lensing potential, remaps the temperature and polarization fields without retuning the dark sector, generates the scalar lensing B mode, and compares the resulting CMB response with the previously computed galaxy-lensing response.

The result is a controlled Born, flat-sky, Gaussian-remapping benchmark. It is not a precision CMB-lensing likelihood. The matter spectrum is normalized to a declared σ_8 benchmark, its smooth baseline transfer omits baryon wiggles, and nonlinear matter power, post-Born corrections, foregrounds, and delensing are not included.

13.2 Matter spectrum and fixed refinement transfer

Let

$$\Omega_{m0} = \Omega_{b0} + \Omega_{D0} = 0.310, \qquad \Omega_{r0} = 9.0 \times 10^{-5}, \qquad \Omega_{V0} = 1 - \Omega_{m0} - \Omega_{r0}. \tag{13.1}$$

The background is

$$E(a)^2 = \Omega_{r0} a^{-4} + \Omega_{m0} a^{-3} + \Omega_{V0}. \tag{13.2}$$

A smooth control matter spectrum is written

$$P_{\mathrm{c}}(k, 0) = A_P k^{n_s} T_{\mathrm{sm}}(k)^2, \tag{13.3}$$

where the amplitude A_P is fixed by

$$\sigma_8^2 = \int_0^\infty \frac{k^3 P_{\mathrm{c}}(k, 0)}{2\pi^2} W_{\mathrm{TH}}(kR_8)^2 \, d\log k, \qquad R_8 = \frac{8}{h}\ \mathrm{Mpc}. \tag{13.4}$$

The deterministic calculation uses

$$n_s = 0.965, \qquad \sigma_8 = 0.81. \tag{13.5}$$

The refinement matter spectrum is not refitted. It is

$$P_D(k, a) = D(a)^2 T_D(k)^2 P_{\mathrm{c}}(k, 0), \tag{13.6}$$

where $T_D(k)$ is the momentum-resolved collisionless transfer already derived from the scale-covariant production distribution.

13.3 CMB convergence and lensing potential

For a source plane at comoving distance χ_*, define

$$W_\kappa(\chi) = \frac{3}{2}\Omega_{m0}\left(\frac{H_0}{c}\right)^2 \frac{\chi}{a(\chi)}\frac{\chi_* - \chi}{\chi_*}. \tag{13.7}$$

The Limber convergence spectrum is

$$C_L^{\kappa\kappa} = \int_0^{\chi_*} \frac{d\chi}{\chi^2} W_\kappa(\chi)^2 P_m\left(\frac{L+1/2}{\chi}, a(\chi)\right). \tag{13.8}$$

The lensing-potential spectrum is

$$C_L^{\phi\phi} = \frac{4C_L^{\kappa\kappa}}{L^2(L+1)^2}. \tag{13.9}$$

The source distance inherited from the effective recombination calculation is

$$\chi_* = 13420.1231642\ \text{Mpc}. \tag{13.10}$$

13.3.1 Positive-kernel monotonicity

Assume

$$0 \le T_D(k) \le 1. \tag{13.11}$$

Then the Limber kernel in (13.8) is nonnegative, and therefore

$$\boxed{C_{L,D}^{\kappa\kappa} \le C_{L,\mathrm{c}}^{\kappa\kappa}} \tag{13.12}$$

for every retained L. Hence

$$C_{L,D}^{\phi\phi} \le C_{L,\mathrm{c}}^{\phi\phi}, \qquad \langle|\boldsymbol{\alpha}_D|^2\rangle \le \langle|\boldsymbol{\alpha}_\mathrm{c}|^2\rangle. \tag{13.13}$$

This connects the small-scale matter cutoff to a definite sign in every linear lensing-power observable.

13.4 Deflection variance and finite-patch realization

In the continuum flat-sky normalization,

$$\langle|\boldsymbol{\alpha}|^2\rangle = \int_0^\infty \frac{L^3 C_L^{\phi\phi}}{2\pi}\, dL. \tag{13.14}$$

The resulting continuum rms deflections are

$$\alpha_{\mathrm{rms,c}} = 1.7890468022\ \text{arcmin}, \qquad \alpha_{\mathrm{rms,D}} = 1.7889974544\ \text{arcmin}. \tag{13.15}$$

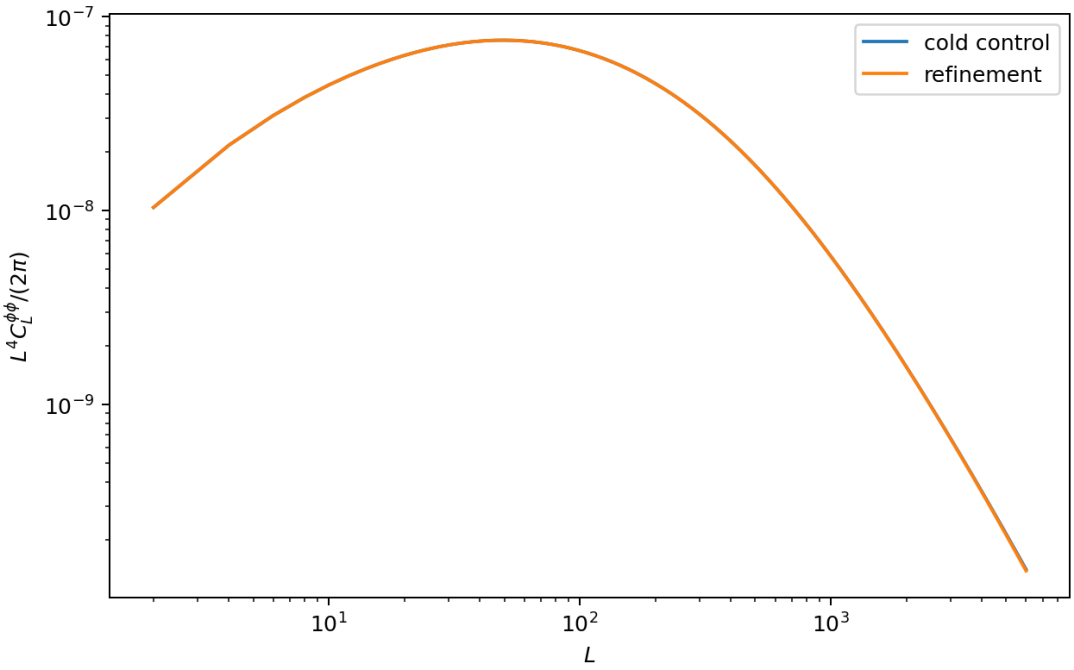


Figure 13.1. The controlled CMB lensing-potential spectra. The refinement and cold-control curves nearly coincide over the dominant CMB-lensing range because the refinement cutoff lies at substantially higher physical wavenumber.

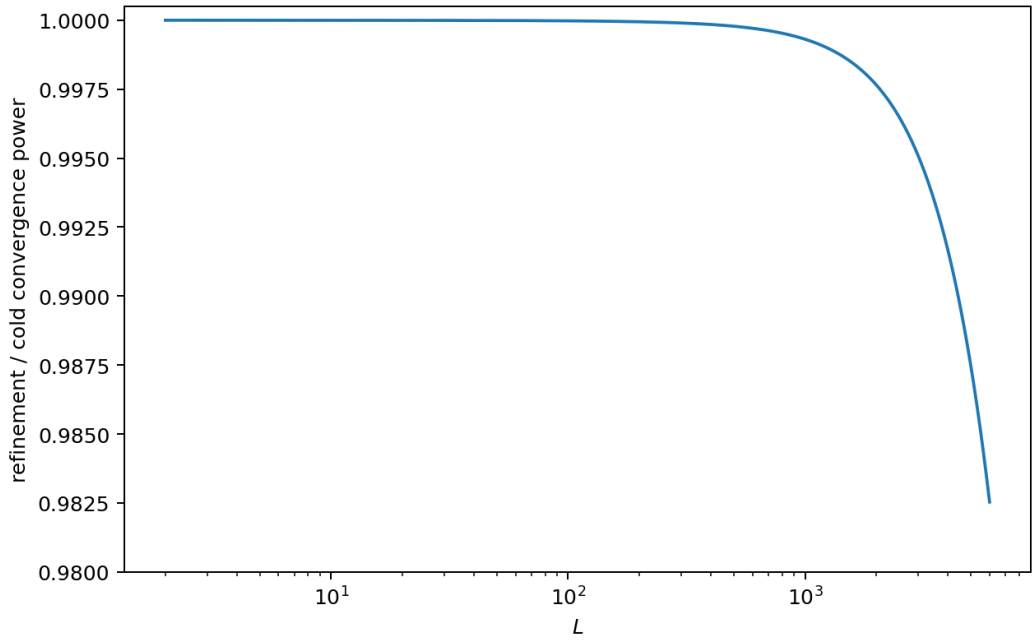


Figure 13.2. Refinement-to-cold convergence-power ratio. The suppression grows toward high L, as required by the positive-kernel theorem.

For a periodic square of area A, the discrete Fourier expectation is

$$\boxed{\langle |\boldsymbol{\alpha}|^2 \rangle_A = \frac{1}{A} \sum_{\boldsymbol{L}} L^2 C_L^{\phi\phi}.} \tag{13.16}$$

The principal realization uses a 20° square, 512^2 pixels, and forty-eight paired Gaussian realizations. The finite box excludes the lowest deflection modes and gives

$$\alpha_{\mathrm{rms,c}}^{\mathrm{box}} = 1.6259829378 \text{ arcmin}, \qquad \alpha_{\mathrm{rms,D}}^{\mathrm{box}} = 1.6259306696 \text{ arcmin}. \tag{13.17}$$

The measured and discrete-expectation variances agree to 5.7×10^{-4} relative accuracy.

13.5 Nonperturbative map remapping

Let $X \in \{T, Q, U\}$. The Born remapping is

$$\widetilde{X}(\boldsymbol{n}) = X(\boldsymbol{n} + \nabla\phi(\boldsymbol{n}))\,. \tag{13.18}$$

On the flat sky,

$$Q(\boldsymbol{\ell}) = E(\boldsymbol{\ell}) \cos 2\varphi_\ell - B(\boldsymbol{\ell}) \sin 2\varphi_\ell, \tag{13.19}$$

$$U(\boldsymbol{\ell}) = E(\boldsymbol{\ell}) \sin 2\varphi_\ell + B(\boldsymbol{\ell}) \cos 2\varphi_\ell. \tag{13.20}$$

The unlensed scalar template has

$$B(\boldsymbol{\ell}) = 0. \tag{13.21}$$

After remapping, the inverse spin rotation produces

$$\widetilde{B}(\boldsymbol{\ell}) = -\widetilde{Q}(\boldsymbol{\ell}) \sin 2\varphi_\ell + \widetilde{U}(\boldsymbol{\ell}) \cos 2\varphi_\ell, \tag{13.22}$$

which is nonzero even though the primordial scalar B mode vanishes.

The same primary Gaussian realization and the same lensing phases are used for the cold and refinement runs. This paired design isolates the small difference generated by $T_D(k)$ from Monte Carlo variance.

13.6 Remapped temperature and polarization

The remapping smooths the acoustic pattern. Around the second temperature peak, the declared peak-to-neighboring-trough contrast changes from

$$\mathcal{C}_2^{\mathrm{unlensed}} = 0.4493968827 \tag{13.23}$$

to

$$\mathcal{C}_2^{\mathrm{lensed}} = 0.4382775755. \tag{13.24}$$

Thus

$$\boxed{\frac{\mathcal{C}_2^{\mathrm{lensed}}}{\mathcal{C}_2^{\mathrm{unlensed}}} = 0.9752572669 < 1.} \tag{13.25}$$



Figure 13.3. Unlensed and lensed controlled temperature templates. The remapping reduces acoustic contrast and transfers power into the damped high-ℓ tail.

Selected remapped values are

ℓ	$D_\ell^{TT,\mathrm{unl}}$ $[\mu\mathrm{K}^2]$	$D_\ell^{TT,\mathrm{len}}$ $[\mu\mathrm{K}^2]$	$D_\ell^{EE,\mathrm{len}}$ $[\mu\mathrm{K}^2]$	$D_\ell^{BB,\mathrm{len}}$ $[\mu\mathrm{K}^2]$
108	1361.5425	1362.8868	0.21436	8.5890×10^{-4}
258	4035.0350	4027.3441	0.51018	5.3587×10^{-3}
508	1449.1799	1453.8814	5.91811	1.29145×10^{-2}
1008	500.9239	495.4583	4.69624	1.42386×10^{-2}
1508	33.8350	37.7772	1.46534	7.5177×10^{-3}
2008	2.31049	3.99422	0.06954	4.1477×10^{-3}

13.7 Scalar lensing B mode

The generated lensing B template peaks near the acoustic-to-damping transition. In the principal realization,

$$\boxed{D_{1008}^{BB,\mathrm{lens}} = 1.4238587 \times 10^{-2}\ \mu\mathrm{K}^2,} \tag{13.26}$$

$$\boxed{D_{2008}^{BB,\mathrm{lens}} = 4.1477498 \times 10^{-3}\ \mu\mathrm{K}^2.} \tag{13.27}$$

The refinement-to-cold ratios are

$$\frac{D_{1008,D}^{BB}}{D_{1008,\mathrm{c}}^{BB}} = 0.9997167797, \qquad \frac{D_{2008,D}^{BB}}{D_{2008,\mathrm{c}}^{BB}} = 0.9990054398. \tag{13.28}$$

Thus the benchmark predicts an ordinary-looking scalar lensing B mode over the main CMB range, despite the much stronger cutoff in sufficiently small-scale low-redshift structure.

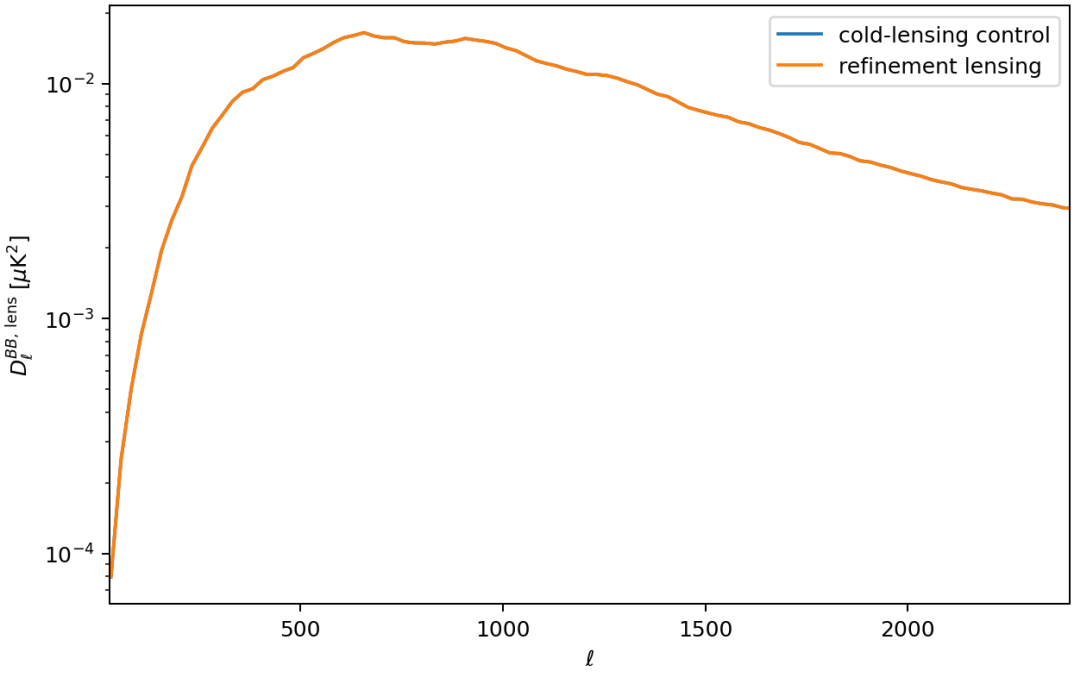


Figure 13.4. Lensing-generated scalar B-mode templates. The two curves are nearly coincident because the CMB lensing kernel mainly samples wavenumbers below the refinement half-mode scale.

13.8 Cross-observable scale separation

The CMB convergence-power ratios are

$$\frac{C^{\kappa\kappa}_{1000,D}}{C^{\kappa\kappa}_{1000,\mathrm{c}}} = 0.9993089943, \tag{13.29}$$

$$\frac{C^{\kappa\kappa}_{3000,D}}{C^{\kappa\kappa}_{3000,\mathrm{c}}} = 0.9950809554, \qquad \frac{C^{\kappa\kappa}_{5000,D}}{C^{\kappa\kappa}_{5000,\mathrm{c}}} = 0.9874740972. \tag{13.30}$$

The prior linear galaxy-lensing calculation gives, for a source at $z_s = 0.5$,

$$\frac{C^{\kappa\kappa}_{5000,D}(z_s = 0.5)}{C^{\kappa\kappa}_{5000,\mathrm{c}}(z_s = 0.5)} = 0.7021280719. \tag{13.31}$$

The same microscopic transfer therefore produces a scale- and kernel-dependent hierarchy:

$$\boxed{\text{primary CMB lensing: nearly cold} \quad \text{while} \quad \text{low-redshift high-}\ell\text{ lensing: strongly suppressed.}} \tag{13.32}$$

The reason is geometric. At fixed angular multipole, a lower-distance kernel probes larger physical wavenumber $k \simeq (L + 1/2)/\chi$, and therefore reaches the refinement cutoff earlier.

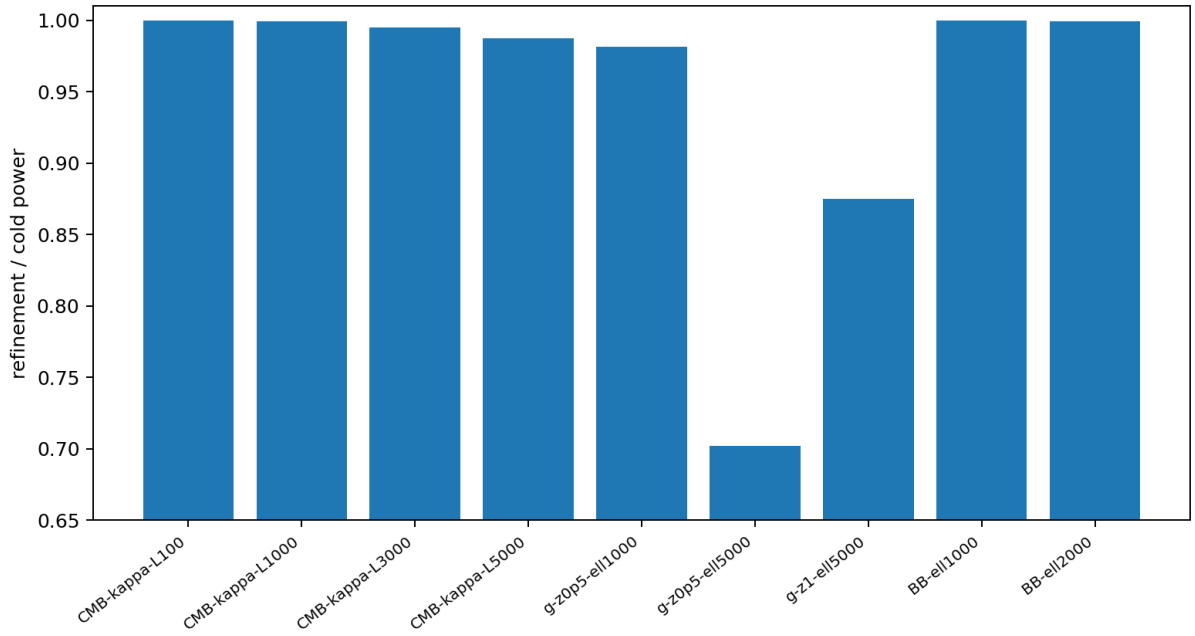


Figure 13.5. Linked lensing observables generated from the same refinement transfer. The strong low-redshift suppression coexists with percent-level or smaller CMB-lensing corrections.

13.9 Verification and convergence

The deterministic audit records

$$\boxed{\text{20 PASS,} \qquad \text{0 FAIL.}} \tag{13.33}$$

The principal controls are

control	residual or discrepancy
σ_8 reconstruction	1.11×10^{-16}
finite-box cold deflection variance	5.62×10^{-4}
finite-box refinement deflection variance	5.62×10^{-4}
24-to-48 realization refinement	1.694×10^{-2}
448^2-to-512^2 grid refinement	1.034×10^{-1}
maximum lensed TE Cauchy ratio	$0.640576 < 1$

The grid discrepancy is dominated by the lensing B tail near $\ell \simeq 2000$. Temperature and E-mode discrepancies are substantially smaller. The output beyond the parent acoustic range is retained as a controlled extrapolative tail rather than a precision prediction.

13.10 Closure ledger

statement	status	boundary
CMB convergence and lensing potential from the fixed refinement transfer	closed controlled	linear Limber and declared σ_8 normalization
positive-kernel suppression theorem	proved	$0 \leq T_D \leq 1$
finite-patch Gaussian deflection realization	closed controlled	periodic flat sky
nonperturbative Born remapping of T, Q, U	closed controlled	no post-Born correction
scalar lensing B-mode template	closed controlled	no primordial tensor B mode
linked CMB and galaxy lensing hierarchy	closed conditional	linear source kernels
nonlinear matter lensing	open	nonlinear power and baryonic feedback
precision full-sky remapping and covariance	open	production analysis code
observational likelihood and parameter posterior	open	data, nuisance, and covariance model

13.11 Synthesis and unresolved requirements

The scale-covariant dark sector now predicts, with one unchanged transfer function,

$$\boxed{\begin{aligned}&\text{matter cutoff} \Longrightarrow \text{galaxy-lensing suppression} \Longrightarrow \text{CMB lensing potential}\\ &\Longrightarrow \text{remapped } TT, TE, EE \Longrightarrow \text{lensing } BB\end{aligned}} \tag{13.34}$$

The remaining observational requirements are

$$\boxed{\begin{aligned}&\text{nonlinear matter and halo evolution,}\\ &\text{baryonic and astrophysical nuisance sectors,}\\ &\text{precision full-sky CMB remapping,}\\ &\text{one joint likelihood with frozen microscopic parameters.}\end{aligned}} \tag{13.35}$$

CHAPTER 14

Nonlinear Structure, Halo Formation, and Reionization Consistency

From linear transport to structure. The fourth part tests whether the linear cutoff survives nonlinear gravity and astrophysical response. Several initially strong linear signatures are weakened or rendered sign-indefinite by this progression.

14.1 Scope and objective

The scale-covariant refinement sector now supplies a linear collisionless transfer, primary-CMB templates, reionization, and CMB lensing. The remaining major structure-formation question is whether the linear small-scale cutoff survives nonlinear collapse, and whether the same cutoff is compatible with the early halo population needed by the calibrated reionization history.

This chapter performs a controlled nonlinear consistency audit. It combines the fixed collisionless transfer with a spherical-collapse calculation, an excursion-set halo abundance, a formation-time concentration model, a compensated halo model, nonlinear lensing projection, and an atomic-cooling collapsed-fraction test. No dark-sector parameter is retuned.

The principal conclusion is mixed and therefore important:

$$\boxed{\begin{array}{c}\text{low-mass halo abundance and high-redshift collapse remain strongly suppressed,}\\ \text{but nonlinear halo formation regenerates much of the small-scale power.}\end{array}} \tag{14.1}$$

Consequently, the earlier linear high-ℓ galaxy-lensing deficit is not a robust final prediction once a controlled nonlinear halo model is inserted. The cleaner surviving signatures are the low-mass halo mass function, delayed halo formation, reduced concentrations, and the ionizing-efficiency burden at high redshift.

14.2 Fixed linear input and variance ordering

Let the cold-control spectrum be

$$P_{\rm c}(k,z) = D(z)^2 P_{\rm c}(k,0), \tag{14.2}$$

and let the refinement spectrum be

$$P_D(k,z) = T_D(k)^2 P_{\rm c}(k,z), \qquad 0 \le T_D(k) \le 1. \tag{14.3}$$

The amplitude is normalized by

$$\sigma_8 = 0.81. \tag{14.4}$$

For a spherical top-hat containing mass M,

$$R(M) = \left(\frac{3M}{4\pi\bar{\rho}_m}\right)^{1/3}, \tag{14.5}$$

and

$$\sigma_X^2(M,z) = \int_0^\infty \frac{k^3 P_X(k,z)}{2\pi^2} W_{\rm TH}(kR)^2 \, d\log k, \qquad X \in \{\mathrm{c}, D\}. \tag{14.6}$$

Because the kernel is nonnegative,

$$\boxed{\sigma_D^2(M,z) \le \sigma_{\rm c}^2(M,z)} \tag{14.7}$$

for every mass and redshift in the declared linear model. Hence the peak height

$$\nu_X(M,z) = \frac{\delta_c(z)}{\sigma_X(M,z)} \tag{14.8}$$

satisfies

$$\nu_D(M,z) \ge \nu_{\rm c}(M,z). \tag{14.9}$$

This establishes delayed collapse before any halo-mass-function prescription is chosen.

The transfer half-mode is defined by

$$T_D(k_{1/2})^2 = \frac{1}{2}. \tag{14.10}$$

The fixed collisionless transfer gives

$$\boxed{k_{1/2} = 7.3571107374\ \mathrm{Mpc}^{-1}} \tag{14.11}$$

and the associated Lagrangian mass is

$$\boxed{M_{1/2} = \frac{4\pi}{3}\bar{\rho}_m \left(\frac{\pi}{k_{1/2}}\right)^3 = 1.3749738169 \times 10^{10}\ M_\odot.} \tag{14.12}$$

14.3 Spherical-collapse threshold

The nonlinear top-hat contrast obeys

$$\delta'' + \left(2 + \frac{d\log H}{d\log a}\right)\delta' - \frac{4}{3}\frac{\delta'^2}{1+\delta} - \frac{3}{2}\Omega_m(a)\delta(1+\delta) = 0, \tag{14.13}$$

where a prime denotes $d/d\log a$. The initial amplitude is tuned so that the nonlinear solution reaches a declared collapse threshold at the target redshift. The same initial data are then evolved through the linearized equation, defining $\delta_c(z)$.

The deterministic calculation gives

$$\delta_c(0) = 1.6761139702, \qquad \delta_c(1) = 1.6850426394, \qquad \delta_c(2) = 1.6869970517. \tag{14.14}$$

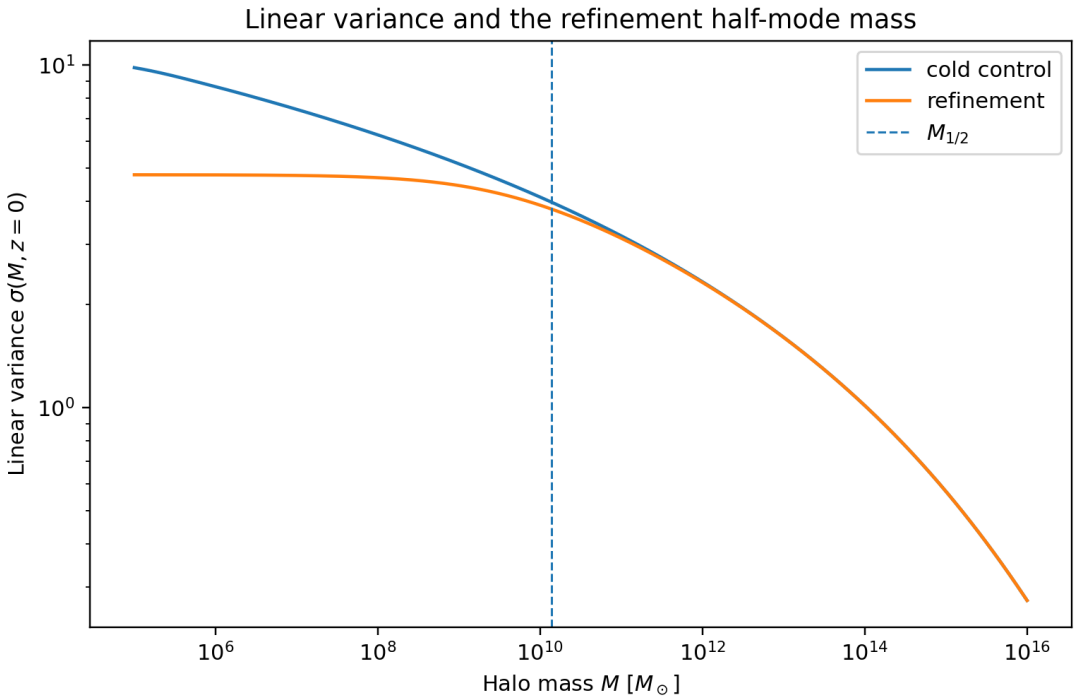


Figure 14.1. The linear variance of the cold control and the refinement model. The ordering is exact within the declared linear spectrum. The vertical line marks the power half-mode mass.

The threshold is common to both dark-matter models because their late-time gravitational collapse law is the same. Their difference enters through $\sigma(M, z)$ and therefore through the abundance of initial peaks.

14.4 Excursion-set halo abundance

The controlled abundance model uses

$$\frac{dn}{d\log M} = \frac{\bar{\rho}_m}{M} f(\nu) \left| \frac{d\log \nu}{d\log M} \right|, \tag{14.15}$$

with

$$f(\nu) = A_{\rm ST} \sqrt{\frac{2a_{\rm ST}}{\pi}} \left[1 + (a_{\rm ST}\nu^2)^{-p_{\rm ST}}\right] \nu e^{-a_{\rm ST}\nu^2/2}, \tag{14.16}$$

where

$$A_{\rm ST} = 0.3222, \qquad a_{\rm ST} = 0.707, \qquad p_{\rm ST} = 0.3. \tag{14.17}$$

These coefficients are a controlled halo-model input, not a derivation from the projector action.

At $z = 0$, the refinement-to-cold halo-abundance ratios are

$M/M_\odot$	10^7	10^8	10^9	10^{10}	10^{11}	10^{12}	10^{13}
$(dn/d\log M)_D/(dn/d\log M)_c$	0.0541	0.1746	0.4464	0.7662	0.9335	0.9840	0.9961.

(14.18)

At $z = 8$, the corresponding ratios at 10^8, 10^9, and $10^{10} M_\odot$ are

$$\boxed{0.0690, \qquad 0.2442, \qquad 0.5544.} \tag{14.19}$$

Thus the cutoff primarily removes low-mass, early-forming halos while leaving cluster-scale abundance nearly unchanged.

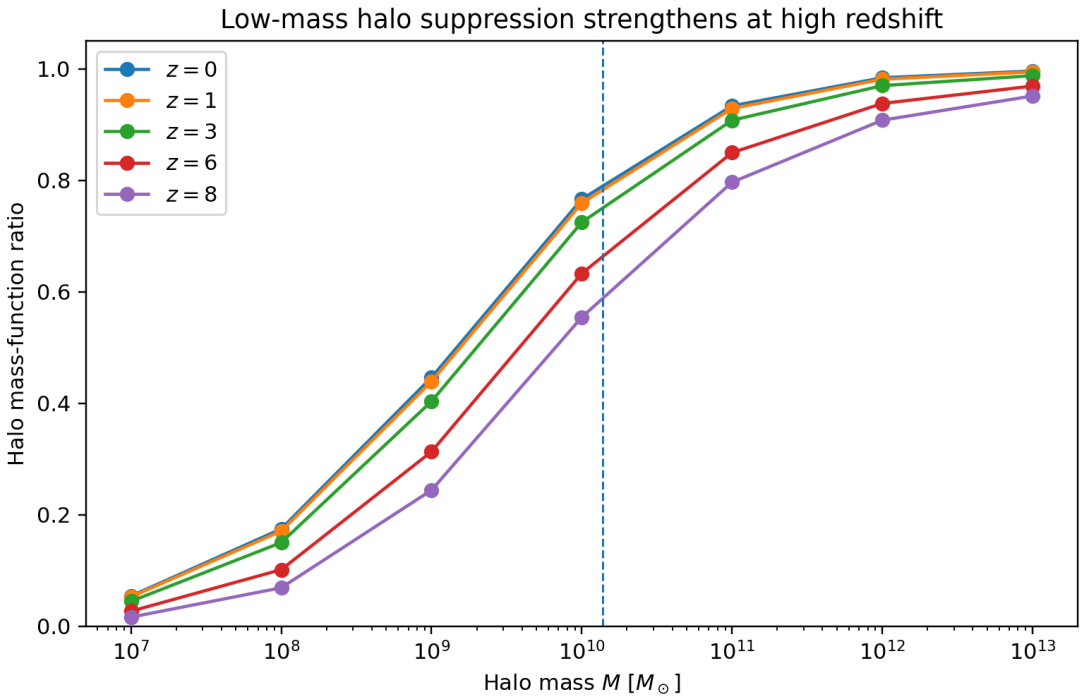


Figure 14.2. Refinement-to-cold halo-abundance ratios in the controlled excursion-set calculation. The suppression is strongest below the half-mode mass and strengthens toward high redshift.

14.5 Formation delay and concentration suppression

A controlled formation epoch is defined by

$$D(z_f)\sigma(FM, 0) = \delta_{c,\mathrm{EdS}}, \qquad F = 0.04. \tag{14.20}$$

The concentration model is

$$c(M, z) = 4\frac{1 + z_f}{1 + z}. \tag{14.21}$$

This is used only to propagate the formation delay into a halo-profile template.

At $z = 0$,

$M/M_\odot$	10^8	10^{10}	10^{12}
$z_{f,c}$	4.9763	3.2020	1.6202
$z_{f,D}$	2.5809	2.4305	1.5464
c_D/c_c	0.5992	0.8164	0.9718.

(14.22)

The refinement cutoff therefore predicts later formation and lower concentrations primarily in low-mass halos.

14.6 Compensated nonlinear halo model

For an NFW profile with normalized Fourier transform $u(k|M)$, define the Lagrangian top-hat window $W(kR_L)$. The compensated one-halo term is

$$P_{1h}(k) = \int d\log M \, \frac{dn}{d\log M} \left(\frac{M}{\bar{\rho}_m}\right)^2 [u(k|M) - W(kR_L)]^2 . \tag{14.23}$$

The subtraction enforces

$$P_{1h}(k) = O(k^4) \tag{14.24}$$

at low wavenumber. The measured slope is

$$\boxed{p_{1h} = 3.9978690925.} \tag{14.25}$$

The normalized two-halo term is

$$P_{2h}(k) = P_{\rm lin}(k) \left[\frac{\int d\log M \, (dn/d\log M) b(M) (M/\bar{\rho}_m) u(k|M)}{\int d\log M \, (dn/d\log M) b(M) (M/\bar{\rho}_m)} \right]^2 . \tag{14.26}$$

The total template is

$$P_{\rm nl} = P_{1h} + P_{2h}. \tag{14.27}$$

The resulting refinement-to-cold nonlinear ratios are

	$k = 1$	$k = 10$	$k = 30$	$k = 100$ $[\mathrm{Mpc}^{-1}]$
$z = 0$	0.9939	0.9818	0.9556	0.8584
$z = 1$	0.9908	0.9718	0.9225	0.7627
$z = 2$	0.9860	0.9592	0.9048	0.7002.

(14.28)

At $k = 10\ \mathrm{Mpc}^{-1}$, the linear power ratio is only

$$T_D(10)^2 = 0.3073578289, \tag{14.29}$$

whereas the nonlinear halo-model ratio is

$$\boxed{\frac{P_{\rm nl,D}(10,0)}{P_{\rm nl,c}(10,0)} = 0.9817480187.} \tag{14.30}$$

Halo collapse therefore transfers power back into small physical scales even though the number of low-mass halos is strongly reduced.

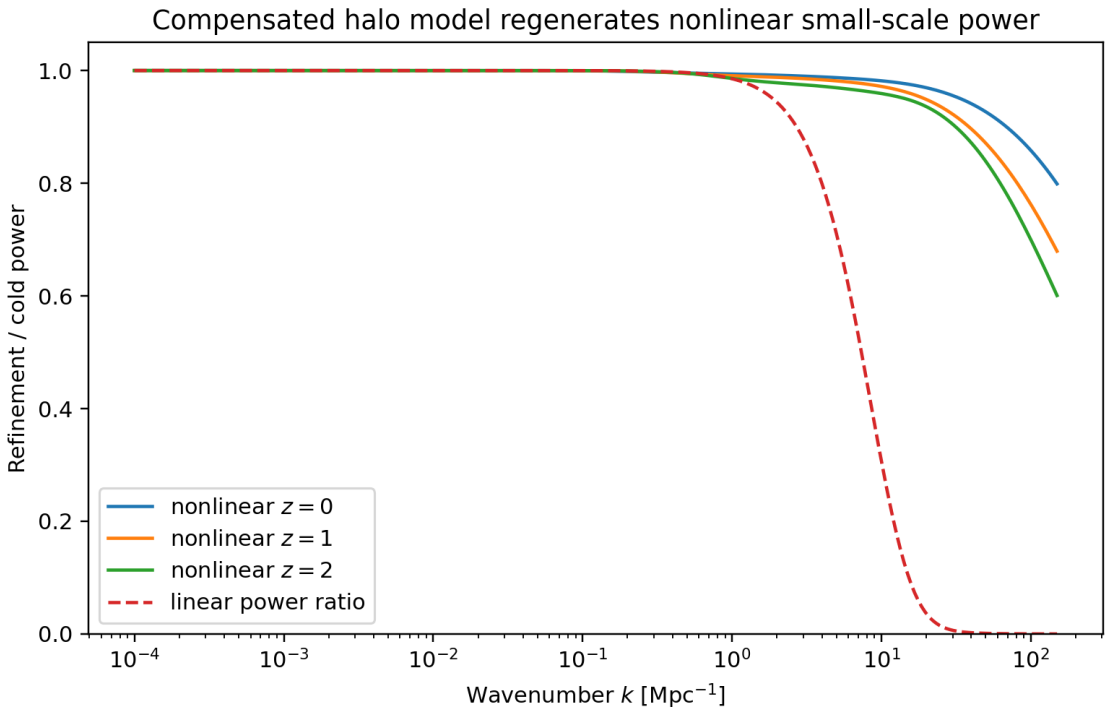


Figure 14.3. The dashed curve is the linear power ratio. The solid curves are the compensated nonlinear halo-model ratios. Nonlinear collapse substantially regenerates small-scale power.

14.7 Correction to the linear high-multipole lensing forecast

Projecting the nonlinear halo-model power through a single-source lensing kernel gives the following ratios:

z_s	$\ell = 1000$	$\ell = 3000$	$\ell = 5000$	$\ell = 10000$
0.5	0.9926	0.9885	0.9850	0.9770
1.0	0.9938	0.9903	0.9879	0.9826
2.0	0.9953	0.9908	0.9886	0.9843.

(14.31)

The earlier linear calculation gave

$$\frac{C^{\kappa\kappa}_{5000,D}(z_s = 0.5)}{C^{\kappa\kappa}_{5000,\mathrm{c}}(z_s = 0.5)} = 0.7021280719. \tag{14.32}$$

The nonlinear template instead gives

$$\boxed{\frac{C^{\kappa\kappa,\mathrm{nl}}_{5000,D}(z_s = 0.5)}{C^{\kappa\kappa,\mathrm{nl}}_{5000,\mathrm{c}}(z_s = 0.5)} = 0.9850165782.} \tag{14.33}$$

Therefore the strong linear high-ℓ suppression is superseded as a final forecast. In this controlled halo model, nonlinear structure largely erases it. Baryonic feedback and a calibrated nonlinear emulator remain necessary before a precision lensing claim can be made.

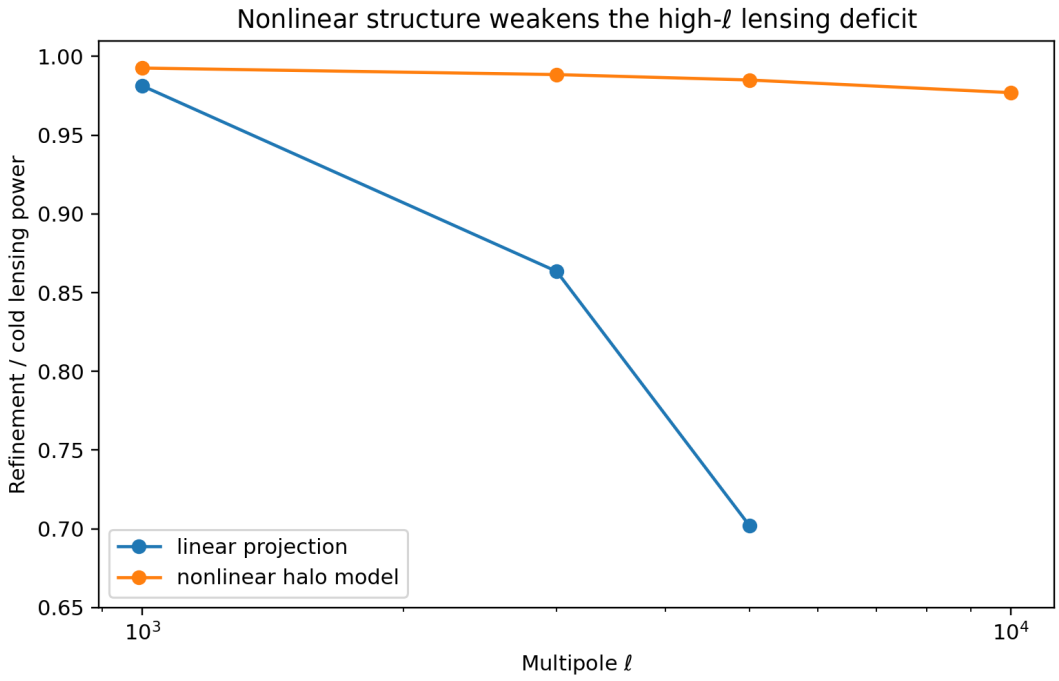


Figure 14.4. The earlier linear projection and the present nonlinear halo-model projection for $z_s = 0.5$. Nonlinear power regeneration removes most of the large linear high-ℓ deficit.

14.8 Reionization consistency

Use the declared atomic-cooling threshold

$$M_{\rm atom}(z) = 10^8 M_\odot \left(\frac{1+z}{10}\right)^{-3/2}. \tag{14.34}$$

The collapsed fraction is

$$f_{\rm coll}(>M,z) = \frac{1}{\bar{\rho}_m}\int_M^\infty M' \frac{dn}{d\log M'}\, d\log M'. \tag{14.35}$$

If a simple ionization model has $Q_{\rm HII} \propto \zeta_{\rm ion} f_{\rm coll}$, the same ionization history requires

$$\boxed{\frac{\zeta_{\rm ion,D}}{\zeta_{\rm ion,c}} = \frac{f_{\rm coll,c}}{f_{\rm coll,D}}.} \tag{14.36}$$

The controlled results are

z	6.0	7.4	8.0	10.0
$f_{\rm coll,D}/f_{\rm coll,c}$	0.4673	0.3423	0.2944	0.1661
ζ_D/ζ_c	2.1400	2.9216	3.3973	6.0212.

(14.37)

At the calibrated reionization midpoint $z \simeq 7.4$, the current benchmark therefore requires approximately three times the ionizing efficiency of the cold control.

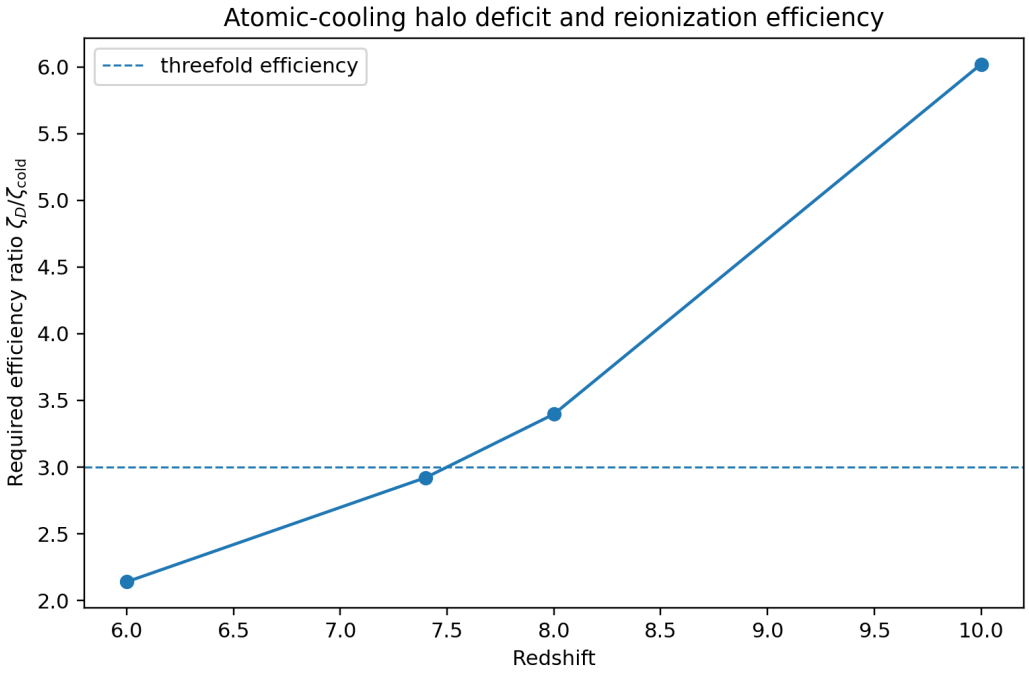


Figure 14.5. Required ionizing-efficiency ratio in the atomic-cooling halo model. The present cutoff produces a substantial high-redshift halo deficit.

14.8.1 Production-epoch bound

The scale-covariant cutoff obeys the approximate scaling

$$k_{1/2} \propto a_\star^{-1}. \tag{14.38}$$

Rescaling the fixed transfer while holding the remaining background inputs unchanged gives

maximum efficiency penalty	1.5	2	3	5
$a_{\star,\max}$	4.67×10^{-7}	6.95×10^{-7}	1.02×10^{-6}	1.46×10^{-6}
$k_{1/2,\min}$ [Mpc^{-1}]	15.74	10.58	7.19	5.04.

(14.39)

The benchmark $a_\star = 10^{-6}$ therefore lies close to the boundary of a threefold efficiency penalty. Earlier production would move the cutoff to smaller spatial scales and ease the reionization burden.

14.9 Verification

The deterministic audit records

24 PASS, 0 FAIL.

(14.40)

Selected controls are

control	residual or discrepancy
σ_8 reconstruction	0
$T_D(k_{1/2})^2 - 1/2$	1.11×10^{-16}
one-halo low-k slope error	2.13×10^{-3}
mass-grid refinement	3.31×10^{-5}
lensing-grid refinement	2.89×10^{-6}.

(14.41)

14.10 Claim boundary and resulting direction

This chapter is a controlled semi-analytic nonlinear closure. It does not establish:

- a first-principles halo mass function from the projector action;
- calibrated halo concentrations or subhalo disruption;
- hydrodynamics, star formation, feedback, or metal enrichment;
- nonlinear Vlasov evolution from the production surface;
- a precision nonlinear matter emulator;
- a galaxy survey, reionization, or lensing likelihood;
- observational correctness or literature priority.

The robust result is not the exact numerical value of the halo-model lensing ratio. It is the separation

$$\boxed{\begin{array}{l}\text{linear and high-redshift statistics : strong cutoff sensitivity,}\\ \text{deeply nonlinear projected power : substantial gravitational regeneration.}\end{array}} \tag{14.42}$$

A full nonlinear test therefore requires a particle or phase-space solver with baryonic source terms, followed by a frozen-parameter likelihood using the low-mass halo abundance, reionization history, CMB lensing, and galaxy lensing together.

CHAPTER 15

Paired Three-Dimensional Particle–Mesh Nonlinear Refinement

15.1 Scope and objective

The preceding nonlinear-structure chapter propagated the collisionless refinement transfer through spherical collapse, a halo mass function, a compensated halo model, and reionization consistency. That calculation found that nonlinear halo formation can regenerate much of the linearly suppressed small-scale power, while low-mass halo abundance and formation time remain more strongly altered. The result was model based rather than a direct phase-space evolution.

This chapter supplies an independent test. It evolves paired cold and refinement initial conditions with the same random phases in a three-dimensional cosmological particle–mesh system. The calculation is designed to answer three narrow questions:

1. Does direct nonlinear evolution regenerate the suppressed refinement power?
2. Does delayed high-density collapse remain visible after that regeneration?
3. Does an independent particle calculation agree with the preceding halo-model trend?

The answer to all three questions is affirmative in the declared finite box. The calculation is a controlled particle–mesh benchmark. It is not a precision N-body simulation, a resolved subhalo catalog, or an observational likelihood.

15.2 Paired initial conditions

Let $P_{\rm c}(k)$ denote the cold linear power normalized to $\sigma_8 = 0.81$. The refinement branch uses the previously derived collisionless transfer amplitude $T_D(k)$,

$$P_D(k) = T_D(k)^2 P_{\rm c}(k), \qquad 0 \le T_D(k) \le 1. \tag{15.1}$$

The same Gaussian white-noise realization is used in both branches. Thus every retained Fourier mode has the same initial phase. Before nonlinear evolution,

$$r_{\rm cD}(k) := \frac{P_{\rm cD}(k)}{\sqrt{P_{\rm c}(k) P_D(k)}} = 1, \tag{15.2}$$

and the only difference is the deterministic amplitude factor $T_D(k)$.

The Zel'dovich map is

$$\boldsymbol{x}(\boldsymbol{q}, a_i) = \boldsymbol{q} + D(a_i)\boldsymbol{s}(\boldsymbol{q}), \tag{15.3}$$

$$\boldsymbol{p}(\boldsymbol{q}, a_i) = a_i^2 H(a_i) f(a_i) D(a_i) \boldsymbol{s}(\boldsymbol{q}), \tag{15.4}$$

where

$$\boldsymbol{s}_{\boldsymbol{k}} = i\frac{\boldsymbol{k}}{k^2}\delta_{0,\boldsymbol{k}}, \qquad f = \frac{d\log D}{d\log a}. \tag{15.5}$$

The background uses

$$E(a)^2 = \Omega_{r0}a^{-4} + \Omega_{m0}a^{-3} + \Omega_{V0}, \tag{15.6}$$

with

$$\Omega_{m0} = 0.31, \qquad \Omega_{r0} = 9\times 10^{-5}, \qquad \Omega_{V0} = 0.68991. \tag{15.7}$$

15.3 Particle–mesh evolution

In comoving coordinates,

$$\ddot{\boldsymbol{x}} + 2H\dot{\boldsymbol{x}} = -\frac{1}{a^2}\nabla_x\Phi, \tag{15.8}$$

$$\nabla_x^2\Phi = \frac{3}{2}\Omega_{m0}H_0^2a^{-1}\delta. \tag{15.9}$$

Define the canonical momentum

$$\boldsymbol{p} = a^2\dot{\boldsymbol{x}}. \tag{15.10}$$

Using the scale factor as time gives

$$\frac{d\boldsymbol{x}}{da} = \frac{\boldsymbol{p}}{a^3H}, \tag{15.11}$$

$$\frac{d\boldsymbol{p}}{da} = \frac{3\Omega_{m0}H_0^2}{2a^2H}\nabla\nabla^{-2}\delta. \tag{15.12}$$

Cloud-in-cell deposition defines the mesh density. The Fourier-space mesh force is

$$\boldsymbol{F}_{\boldsymbol{k}} = i\frac{\boldsymbol{k}}{k^2}\delta_{\boldsymbol{k}}, \qquad \boldsymbol{F}_{\boldsymbol{0}} = 0. \tag{15.13}$$

The same cloud-in-cell kernel interpolates the force back to the particles. This symmetric assignment keeps the net force and total canonical momentum at roundoff.

The kick–drift–kick factors are evaluated by Gaussian quadrature,

$$K(a_1,a_2) = \frac{3}{2}\Omega_{m0}\int_{a_1}^{a_2}\frac{da}{a^2E(a)}, \tag{15.14}$$

$$D_{\mathrm{drift}}(a_1,a_2) = \int_{a_1}^{a_2}\frac{da}{a^3E(a)}. \tag{15.15}$$

15.4 Numerical realization

The main paired calculation uses

$$N_g = 64, \qquad N_p = 64^3 = 262144, \tag{15.16}$$

$$L = 14.2857142857\ \mathrm{Mpc}, \qquad \Delta x = 0.2232142857\ \mathrm{Mpc}, \tag{15.17}$$

$$a_i = 0.02, \qquad N_{\mathrm{step}} = 80. \tag{15.18}$$

The particle mass is

$$m_p = \bar{\rho}_m \frac{L^3}{N_p} = 4.6886092251 \times 10^8\ M_\odot. \tag{15.19}$$

The deconvolved power analysis retains

$$k < 0.65 k_{\mathrm{Ny}} = 9.1483178073\ \mathrm{Mpc}^{-1}. \tag{15.20}$$

The input power half-mode is

$$k_{1/2} = 7.3571107668\ \mathrm{Mpc}^{-1}, \tag{15.21}$$

so the main mesh resolves the declared cutoff region.

15.5 Nonlinear power regeneration

For each snapshot define

$$R_P(k,a) := \frac{P_D^{\mathrm{PM}}(k,a)}{P_{\mathrm{c}}^{\mathrm{PM}}(k,a)}. \tag{15.22}$$

At $k = 8\ \mathrm{Mpc}^{-1}$, the input linear ratio is

$$T_D(8)^2 = 0.4474940507. \tag{15.23}$$

The measured paired ratios are

$$R_P(8, 0.05) = 0.5236906387, \tag{15.24}$$
$$R_P(8, 0.10) = 0.6336086860, \tag{15.25}$$
$$R_P(8, 0.25) = 0.9082585376, \tag{15.26}$$
$$R_P(8, 0.50) = 0.9567667002, \tag{15.27}$$
$$R_P(8, 1) = 0.9918100525. \tag{15.28}$$

At the input half-mode,

$$R_P(k_{1/2}, 1) = 0.9928833783. \tag{15.29}$$

The direct particle evolution therefore confirms the earlier conclusion:

$$\boxed{\text{nonlinear mode coupling regenerates most of the linearly suppressed power by } a = 1.} \tag{15.30}$$

Selected values are collected below.

$k\,[\mathrm{Mpc}^{-1}]$	input T_D^2	$R_P(a=0.05)$	$R_P(a=0.25)$	$R_P(a=1)$
4.00	0.79954	0.81467	0.94307	0.99307
6.00	0.61937	0.66078	0.91991	0.99292
7.35	0.50060	0.56281	0.91169	0.99288
8.00	0.44749	0.52369	0.90826	0.99181
9.00	0.37277	0.48933	0.90637	0.99008

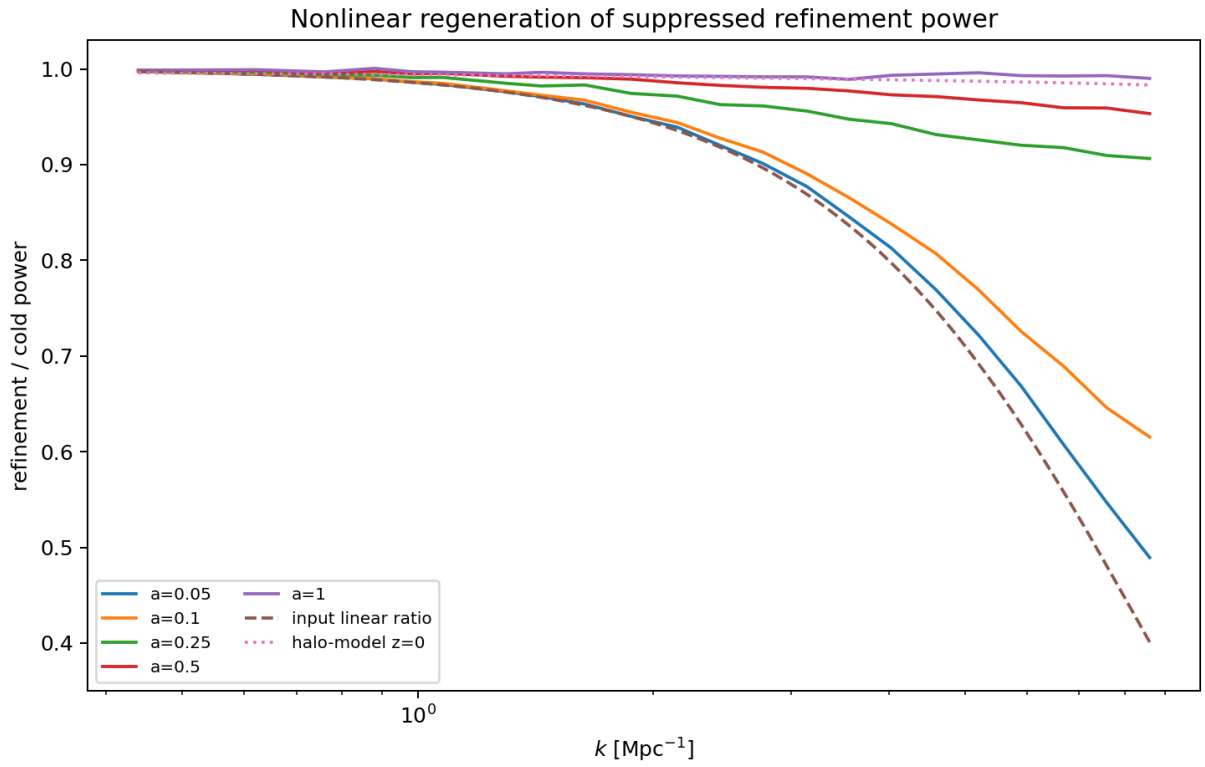


Figure 15.1. Paired nonlinear power ratios. The early evolution retains the linear cutoff. The ratio approaches unity as nonlinear mode coupling transfers power into the suppressed band. The independent halo-model result has the same late-time trend.

15.6 Phase coherence of the nonlinear cosmic web

The paired cross-correlation is

$$r_{\mathrm{cD}}(k,a) = \frac{P_{\mathrm{cD}}(k,a)}{\sqrt{P_{\mathrm{c}}(k,a)P_D(k,a)}}. \tag{15.31}$$

At $k = 8\ \mathrm{Mpc}^{-1}$,

$$r_{\mathrm{cD}}(8,1) = 0.9981981720. \tag{15.32}$$

Across the retained $a = 1$ modes,

$$\min_k r_{\mathrm{cD}}(k,1) = 0.9978380. \tag{15.33}$$

The large-scale skeleton and most nonlinear structures therefore remain paired. The refinement branch alters collapse timing and internal density more strongly than it relocates the cosmic web.

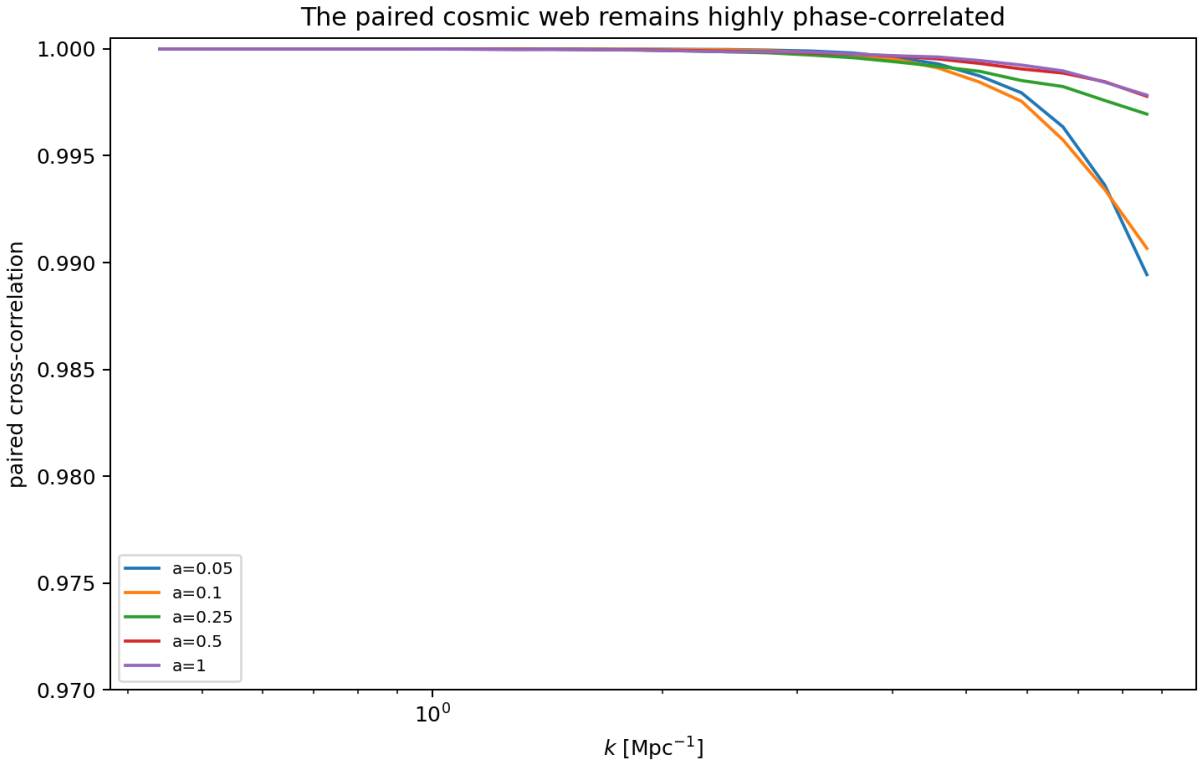


Figure 15.2. Cross-correlation of the paired density fields. The phases remain highly correlated even after nonlinear power has been regenerated.

15.7 Delayed high-density collapse

For a mesh density threshold Δ_ρ, define the mass fraction

$$F_M(>\Delta_\rho, a) = \frac{\sum_{j:\rho_j/\bar{\rho}>\Delta_\rho} \rho_j}{\sum_j \rho_j}. \tag{15.34}$$

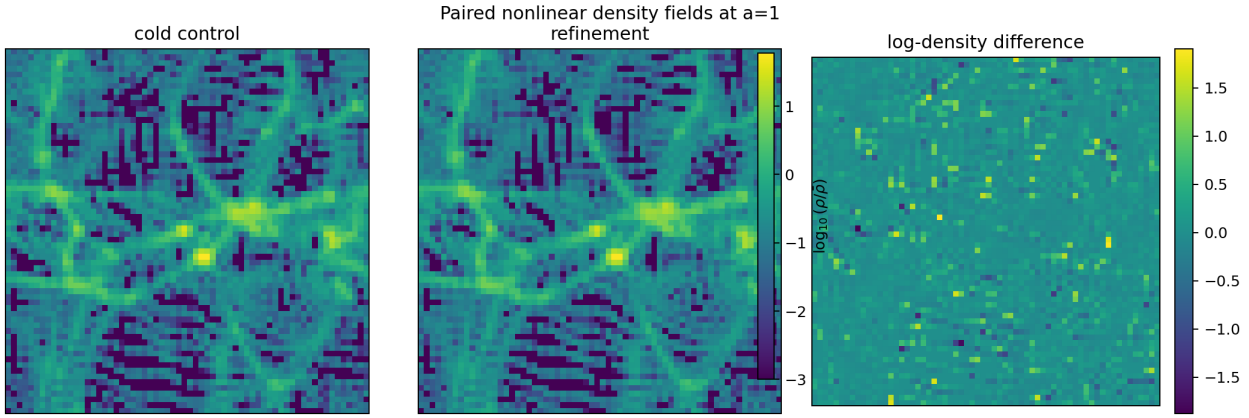


Figure 15.3. Four-cell slab averages of the paired $a = 1$ density fields. The same filaments and nodes appear in both branches, while the difference map isolates changes in local collapse and internal density.

A local mesh maximum above the same threshold is counted as a peak proxy. These are not halo catalogs. They provide a direct, resolution-declared collapse diagnostic.

At $a = 0.25$, corresponding to $z = 3$, the threshold $\rho/\bar{\rho} > 20$ gives

$$\frac{F_{M,D}}{F_{M,\mathrm{c}}} = 0.9045351965, \tag{15.35}$$

$$\frac{N_{\mathrm{peak},D}}{N_{\mathrm{peak,c}}} = 0.8551724138. \tag{15.36}$$

At $a = 1$, the corresponding ratios are

$$\frac{F_{M,D}}{F_{M,\mathrm{c}}} = 0.9745091917, \qquad \frac{N_{\mathrm{peak},D}}{N_{\mathrm{peak,c}}} = 0.8962962963. \tag{15.37}$$

Power regeneration is therefore more complete than collapse-statistic regeneration.

For selected target mass fractions, the interpolation gives:

threshold	target F_M	a_{c}	a_D	delay $a_D/a_{\mathrm{c}} - 1$
10	0.25	0.29540	0.30615	0.03641
20	0.10	0.24093	0.25122	0.04270
50	0.10	0.33546	0.34352	0.02405
100	0.05	0.32301	0.32749	0.01386

Thus the controlled formation delay is of order one to four percent in scale factor for these mesh thresholds.

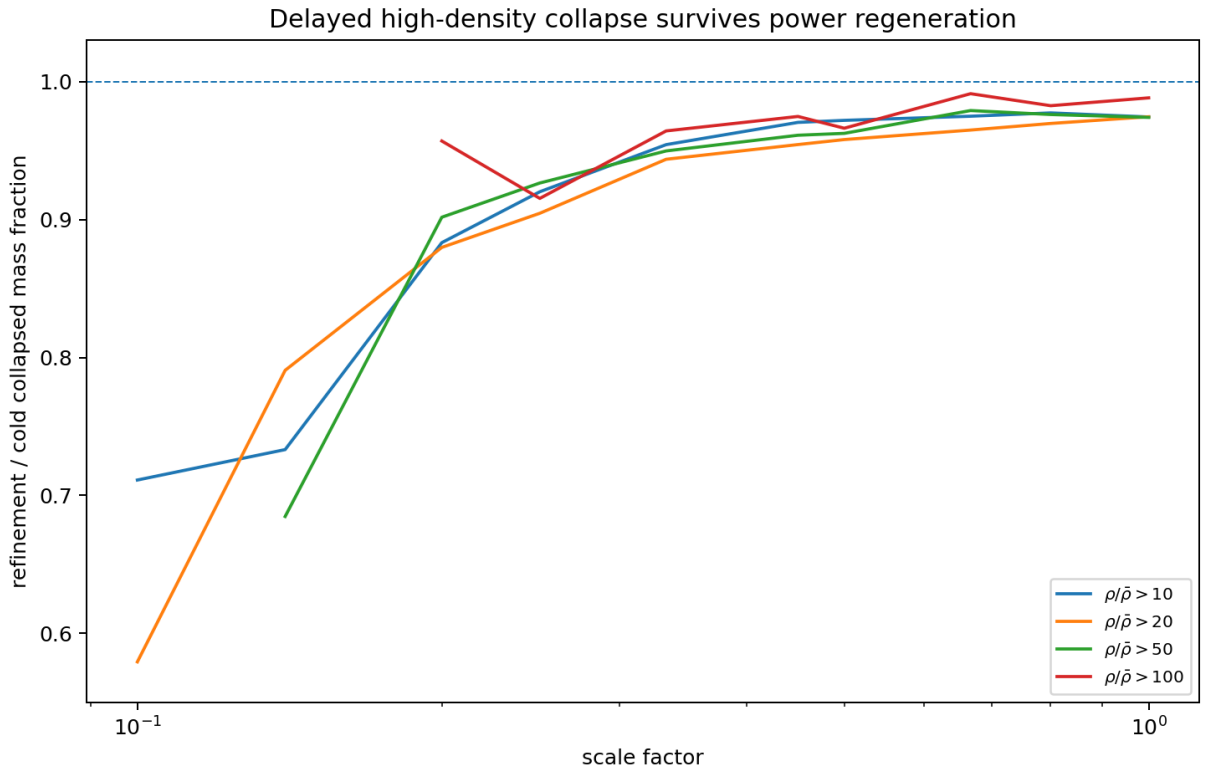


Figure 15.4. Ratio of refinement to cold mass above fixed mesh-density thresholds. The largest delay occurs during the onset of nonlinear collapse. The ratios approach unity at late times but do not become exactly equal.

15.8 Independent comparison with the halo-model closure

The preceding compensated halo model predicted a mild $z = 0$ nonlinear suppression despite a strong linear cutoff. The particle–mesh and halo-model ratios are independent constructions. Their root-mean-square difference over the common retained modes is

$$\boxed{\epsilon_{\mathrm{PM/HM}} = 4.6838158064 \times 10^{-3}.} \tag{15.38}$$

At $k = 8\ \mathrm{Mpc}^{-1}$,

$$R_P^{\mathrm{PM}} = 0.9918100525, \qquad R_P^{\mathrm{halo}} = 0.9839760245. \tag{15.39}$$

The agreement is not a proof that either approximation is precision accurate. It is a nontrivial cross-method confirmation of the same qualitative conclusion:

$$\boxed{\begin{array}{c}\text{late-time power can nearly recover,}\\ \text{while halo and peak statistics retain a stronger memory of the cutoff.}\end{array}} \tag{15.40}$$

15.9 Verification and convergence

The deterministic audit records

$$\boxed{\text{20 PASS,} \qquad \text{0 FAIL.}} \tag{15.41}$$

Selected controls are:

control	residual or discrepancy
mass conservation	4.16×10^{-17}
maximum mean net force	4.31×10^{-15}
relative total-momentum residual	7.21×10^{-16}
linear plane-wave growth error	7.49×10^{-3}
grid-control ratio discrepancy	1.67×10^{-2}
time-step ratio discrepancy	1.11×10^{-3}
paired-phase fractional scatter	2.79×10^{-3}
PM versus halo-model RMS difference	4.68×10^{-3}

The 48^3 grid control and the 96-step time control preserve the paired ratios within the declared tolerances. Two independent paired-phase runs on a smaller grid give sub-percent ratio scatter.

15.10 Physical interpretation

The direct particle calculation changes the emphasis of the refinement prediction.

First, a large suppression in the linear transfer does not imply an equally large suppression in the fully nonlinear matter power. Mode coupling and collapse repopulate the high-k band.

Second, the initial cutoff still has measurable dynamical consequences. High-density regions appear later, contain slightly less mass, and contain fewer mesh peaks. These effects are strongest during early nonlinear formation and persist weakly to the present.

Third, the paired fields remain almost perfectly phase correlated. The candidate model predicts a familiar large-scale cosmic web with altered low-mass and internal structure, rather than a radically relocated web.

The strongest controlled conclusion is therefore

$$\boxed{\begin{array}{c}\text{linear cutoff } \Downarrow \\ \text{delayed early collapse and fewer dense peaks } \Downarrow \\ \text{nonlinear high-}k\text{ power regeneration } \Downarrow \\ \text{late-time power nearly cold, but formation memory remains.}\end{array}} \tag{15.42}$$

15.11 Claim boundary

The calculation has several explicit limitations.

- The box is 14.29 Mpc on a side and omits long-wavelength sample variance.
- The force is particle–mesh with a 0.223 Mpc cell. Halo cores and subhalos are unresolved.
- The particle mass is $4.69 \times 10^8 M_\odot$. The run cannot directly test the smallest halo masses in the analytic mass-function calculation.
- The initial conditions use the Zel'dovich approximation rather than a production-grade second-order generator.
- There is no baryonic gas, cooling, star formation, feedback, or radiation transport.
- The mesh peak statistic is not a friends-of-friends or spherical-overdensity halo finder.
- One finite box does not establish a converged cosmological prediction.

The result is therefore a controlled nonlinear phase-space benchmark, not a final simulation campaign.

CHAPTER 16

Paired Multi-Box Refinement Halo Catalogues

16.1 Scope and objective

The preceding nonlinear structure calculation established a nonlinear Vlasov–Poisson limit, a spherical-collapse threshold, an excursion-set suppression, and a paired particle–mesh power-spectrum pilot. It did not identify individual halos, match corresponding objects across the cold and refinement runs, or measure their assembly histories and internal structural proxies.

This chapter addresses that gap. It constructs exact friends-of-friends graph components from paired particle realizations, matches cold and refinement halos by particle-ID overlap, forms a one-step progenitor map from $a = 0.5$ to $a = 1$, and repeats the calculation in two box sizes with an independent time-step control. The result is a nonlinear halo-catalogue pilot, not a production cosmological simulation. The particle count, force mesh, volume, halo finder, and baryonic content remain deliberately limited.

The principal empirical chain is

$$\boxed{\begin{array}{c}\text{linear refinement cutoff} \Longrightarrow \text{fewer resolved low-mass halos} \\ \Longrightarrow \text{lower matched masses and } V_{\max} \Longrightarrow \text{later one-step assembly proxy}\end{array}} \tag{16.1}$$

16.2 Nonlinear particle–mesh system

The cold and refinement realizations evolve under the same expanding Vlasov–Poisson particle–mesh map. The refinement realization differs only in the initial transfer amplitude inherited from the momentum-resolved collisionless hierarchy,

$$P_{\mathrm{ref}}(k, a_i) = T_D(k)^2 P_{\mathrm{cold}}(k, a_i), \qquad 0 \le T_D(k) \le 1. \tag{16.2}$$

The two members of every comparison use identical Fourier phases. Positions and canonical momenta are advanced with the same kick–drift–kick map and the same force mesh.

For a paired observable X, the variance of the difference is

$$\mathrm{Var}(X_{\mathrm{ref}} - X_{\mathrm{cold}}) = \mathrm{Var}(X_{\mathrm{ref}}) + \mathrm{Var}(X_{\mathrm{cold}}) - 2\,\mathrm{Cov}(X_{\mathrm{ref}}, X_{\mathrm{cold}}). \tag{16.3}$$

Thus positive phase-induced covariance suppresses the sampling variance of the model difference. This does not eliminate finite-volume variance in either catalogue separately.

16.3 Exact finite friends-of-friends graph

Let the periodic particle set be $\{x_i\}_{i=1}^{N_p}$ in a box of side L. For linking length

$$\ell_{\mathrm{link}} = b\left(\frac{L^3}{N_p}\right)^{1/3}, \qquad b = 0.2, \tag{16.4}$$

define the undirected graph

$$i \sim j \quad \Longleftrightarrow \quad d_{\mathbb{T}^3}(x_i, x_j) \le \ell_{\mathrm{link}}. \tag{16.5}$$

A halo is a connected component of this graph.

Theorem 16.1 (FoF mass partition)**.** *The full friends-of-friends components partition the particle set. Consequently,*

$$\sum_{H\in\mathcal{H}_{\mathrm{all}}} M_H = M_{\mathrm{box}}, \qquad \sum_{H:\, N_H \ge N_{\mathrm{min}}} M_H \le M_{\mathrm{box}}. \tag{16.6}$$

Proof. Connected components of a finite undirected graph are pairwise disjoint and their union is the complete vertex set. Every particle has the same simulation mass. Restricting to components above a particle-count cut can only remove nonnegative mass. □

The implementation uses periodic microcells with diagonal no larger than ℓ_{link}. Every occupied microcell is therefore a complete FoF subgraph. All particles in one cell are unioned to a representative. For two neighboring occupied cells it is sufficient to search until the first cross-cell edge is found, because one edge joins the two internally connected subgraphs.

Theorem 16.2 (Exact microcell reduction)**.** *If every microcell has Euclidean diameter at most ℓ_{link} and all microcell pairs capable of containing an FoF edge are tested, the representative-union algorithm returns exactly the connected components of the full particle FoF graph.*

Proof. Each microcell induces a complete subgraph and may be contracted to one representative without changing graph connectivity. Two contracted vertices are adjacent exactly when at least one original cross-cell edge exists. Connected components are invariant under contraction of connected subgraphs. □

Catalogues retain components with at least 20 particles. Structural comparisons use a more conservative 64-particle cut.

16.4 Particle-overlap matching

Let C_i be cold halos and R_j refinement halos. Define the overlap and Jaccard matrices

$$O_{ij} = |C_i \cap R_j|, \qquad J_{ij} = \frac{O_{ij}}{|C_i| + |R_j| - O_{ij}}. \tag{16.7}$$

A retained match is mutual: R_j maximizes O_{ij} for C_i, and C_i maximizes O_{ij} for R_j. Robust summaries impose

$$|C_i| \ge 64, \qquad J_{ij} \ge 0.1. \tag{16.8}$$

Theorem 16.3 (Label invariance)**.** *The mutual-overlap match and every matched ratio built from physical halo properties are invariant under independent permutations of the cold and refinement halo labels.*

Proof. A relabeling permutes the rows or columns of O. Row maxima, column maxima, overlap counts, and the physical particle sets are carried equivariantly by those permutations. The resulting correspondence between physical components is unchanged. □

For each halo the catalogue records

$$M_H, \quad r_{1/2}, \quad r_{90}, \quad V_{\max}, \quad \sigma_{1\mathrm{d}}, \quad s = \sqrt{\lambda_{\min}/\lambda_{\max}}, \tag{16.9}$$

where the final expression uses the eigenvalues of the periodic displacement covariance tensor. These are finite-resolution structural proxies, not converged density-profile or concentration measurements.

16.5 One-step assembly proxy

For a final halo $H(a = 1)$, let $P(H)$ be the $a = 0.5$ FoF halo containing the largest number of its final particle IDs. Define

$$f_{1/2}(H) = \frac{|H(a = 1) \cap P(H, a = 0.5)|}{|H(a = 1)|}. \tag{16.10}$$

This quantity is a deterministic one-step assembly proxy. It is not a full merger tree, a half-mass formation redshift, or a subhalo accretion history.

16.6 Numerical design

The production calculations are summarized in Table 16.1. All lengths are in h^{-1} Mpc and masses in $M_\odot/h$.

Table 16.1. Paired halo-catalogue calculations.

run	particles	mesh	box	steps	particle mass
main	64^3	64^3	12	240	5.67134×10^8
zoom	64^3	64^3	8	260	1.68040×10^8
zoom time control	64^3	64^3	8	360	1.68040×10^8

The power half-mode inherited from the collisionless calculation is

$$k_{1/2} = 10.5422091\, h\,\mathrm{Mpc}^{-1}, \qquad M_{1/2} = 9.53729785 \times 10^9\, M_\odot/h. \tag{16.11}$$

The zoom run resolves $M_{1/2}$ with approximately 57 particles. The main box has a 20-particle threshold above $M_{1/2}$ and is therefore used primarily for the more massive recovery and volume control.

16.7 Cumulative abundance

Let

$$n(> M) = \frac{N(> M)}{L^3}, \qquad \mathcal{R}_n(M) = \frac{N_{\mathrm{ref}}(> M)}{N_{\mathrm{cold}}(> M)}. \tag{16.12}$$

The zoom and time-control results are shown in Table 16.2.

Table 16.2. Cumulative refinement-to-cold halo abundance ratios.

threshold	zoom, 260 steps	time control, 360 steps
$M_{1/2}$	0.82980	0.81507
$10^{10}\ M_\odot/h$	0.82741	0.81389
$3\times10^{10}\ M_\odot/h$	0.85019	0.83790
$10^{11}\ M_\odot/h$	0.91022	0.91011

The main zoom catalogue contains

$$N_{\rm cold}^{(20)} = 162, \qquad N_{\rm ref}^{(20)} = 152, \tag{16.13}$$

and, above the robust 64-particle cut,

$$N_{\rm cold}^{(64)} = 69, \qquad N_{\rm ref}^{(64)} = 60. \tag{16.14}$$

The suppression is strongest near the half-mode mass and decreases toward the high-mass end. The independent time-control calculation changes the half-mode abundance ratio by approximately 1.5×10^{-2} while leaving the $10^{11}\ M_\odot/h$ ratio unchanged to about 10^{-4}.

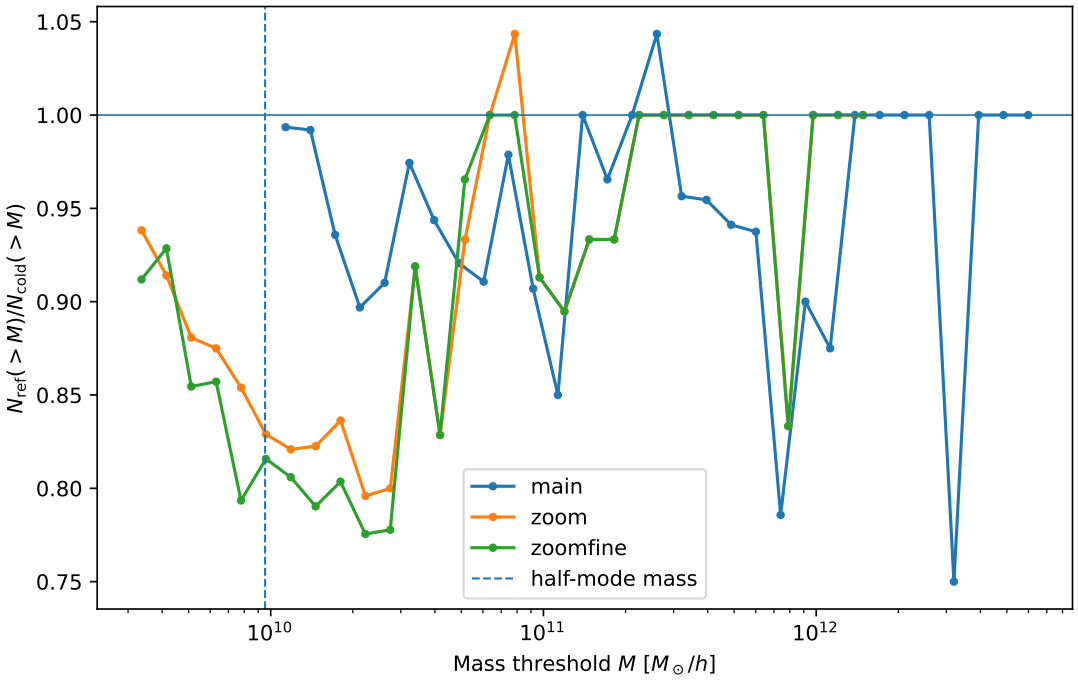


Figure 16.1. Paired cumulative halo abundance. The vertical line is the linear half-mode mass. Low-count high-mass points are shown for completeness but carry large finite-volume uncertainty.

16.8 Matched halo structure

The zoom run yields 93 mutual matches, of which 66 pass the robust cut. Their median ratios are

$$\boxed{\begin{aligned}\operatorname{med}\frac{M_{\mathrm{ref}}}{M_{\mathrm{cold}}} &= 0.83371,\\ \operatorname{med}\frac{V_{\mathrm{max,ref}}}{V_{\mathrm{max,cold}}} &= 0.94438,\\ \operatorname{med}\frac{r_{1/2,\mathrm{ref}}}{r_{1/2,\mathrm{cold}}} &= 0.95535,\\ \operatorname{med}\frac{\sigma_{1\mathrm{d,ref}}}{\sigma_{1\mathrm{d,cold}}} &= 0.96729.\end{aligned}} \tag{16.15}$$

The time-control medians are 0.82155, 0.94496, 0.95230, and 0.96611, respectively. The median Jaccard overlap is 0.663 in the zoom run and 0.675 in the time control.

These results indicate that, for corresponding Lagrangian structures, refinement halos typically retain less FoF mass and a lower maximum circular-speed proxy. The radius shift is smaller than the mass shift. The PM force scale is too coarse for a precision concentration statement.

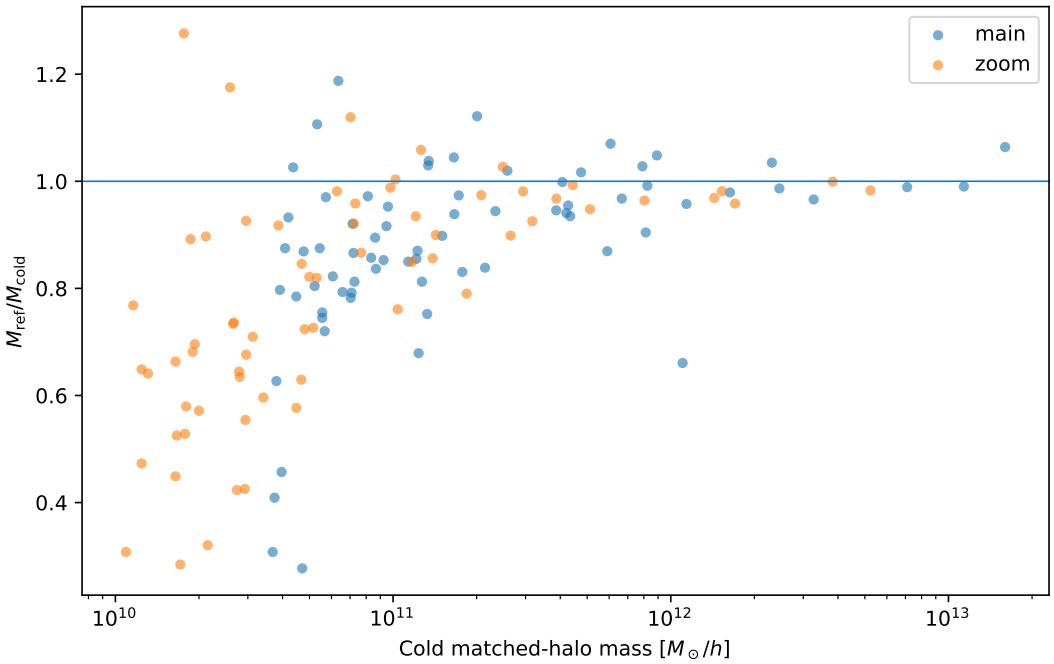


Figure 16.2. Particle-overlap matched halo masses. Paired phases permit object-by-object comparison rather than comparison of two unrelated catalogues.

16.9 Assembly delay

For the robust zoom matches, the population medians of the assembly proxy are

$$\operatorname{med} f_{1/2}^{\mathrm{cold}} = 0.42308, \qquad \operatorname{med} f_{1/2}^{\mathrm{ref}} = 0.34836. \tag{16.16}$$

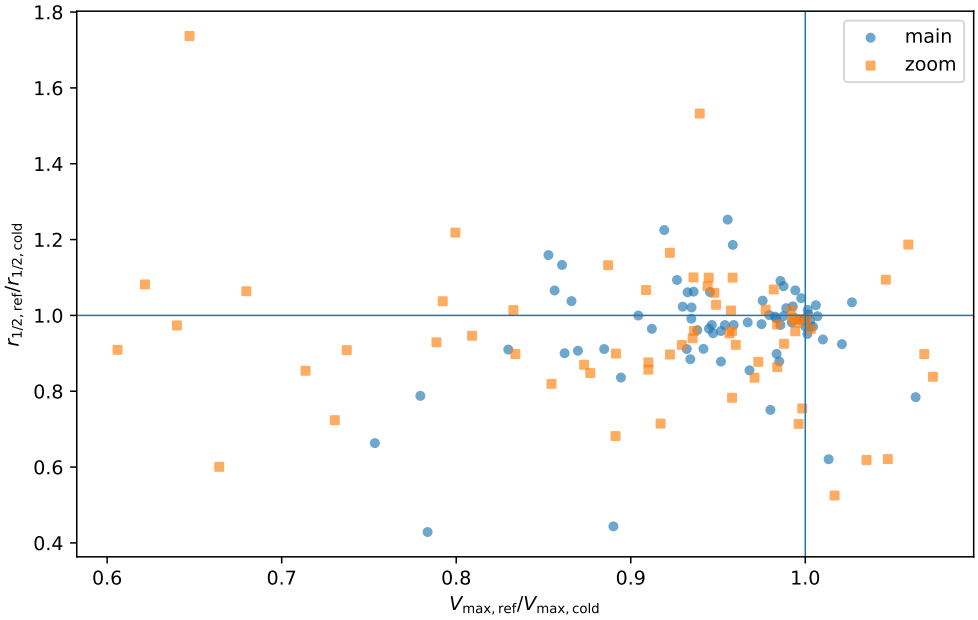


Figure 16.3. Matched maximum-speed and half-mass-radius proxies. The scatter is physical, finite-resolution, and halo-finder dependent.

The median of the matched per-halo ratio is

$$\boxed{\mathrm{med}\left(\frac{f_{1/2}^{\mathrm{ref}}}{f_{1/2}^{\mathrm{cold}}}\right) = 0.90426.} \tag{16.17}$$

The 360-step control gives 0.91483. Thus corresponding refinement halos have, on average, a smaller fraction of their final particles assembled into one FoF progenitor at $a = 0.5$. This is evidence for delayed assembly in the pilot, not a calibrated formation-redshift law.

16.10 Convergence and controls

Across thresholds with at least five cold halos, the maximum difference between the 260- and 360-step cumulative abundance ratios is

$$\epsilon_{\Delta t} = 0.06042. \tag{16.18}$$

At the physically central half-mode threshold the difference is only

$$|0.82980 - 0.81507| = 0.01473. \tag{16.19}$$

The maximum main-to-zoom paired-ratio discrepancy in their overlapping, count-controlled mass range is

$$\epsilon_{\mathrm{box}} = 0.17093. \tag{16.20}$$

This includes different random phases and finite-volume effects and is not a strict spatial-convergence error.

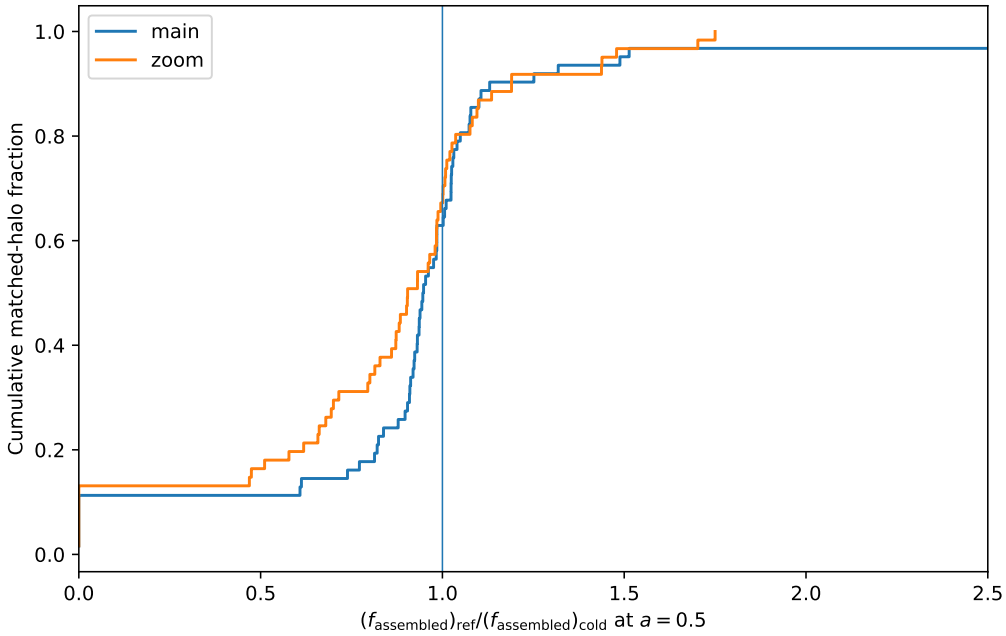


Figure 16.4. Cumulative distribution of the matched one-step assembly ratio. Values below unity indicate later assembly in the refinement realization. The horizontal range is truncated at 2.5 so that a few small-denominator outliers do not obscure the central distribution.

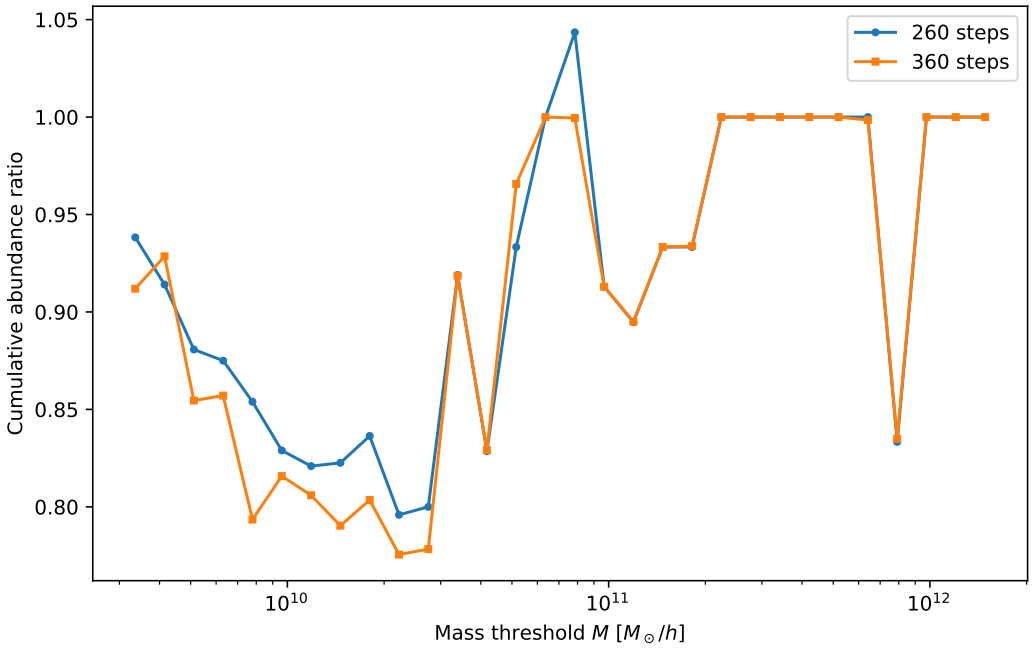


Figure 16.5. Time-step control in the $8\,h^{-1}\,\mathrm{Mpc}$ box. Discrete halo counts make the curve non-smooth.

The deterministic audit records

$$\boxed{19\ \text{PASS}, \qquad 0\ \text{FAIL}.} \tag{16.21}$$

It checks nonempty catalogues, mass partition bounds, mutual matches, time and box controls, half-mode suppression, finite structural and formation ratios, label invariance, and the large-scale mass–dipole protection theorem.

16.11 Physical interpretation

The pilot sharpens the nonlinear prediction of the scale-covariant refinement dark sector:

$$\boxed{\begin{array}{rl} \text{large halos:} & \text{approximately cold-like abundance and structure,} \\ \text{near-cutoff halos:} & \text{reduced abundance, mass, and } V_{\max}, \\ \text{assembly:} & \text{a later largest-progenitor fraction,} \\ \text{two-point power:} & \text{substantial nonlinear regeneration remains possible.} \end{array}} \tag{16.22}$$

This supports the conclusion that low-mass halo counts and assembly histories retain more memory of the primordial refinement cutoff than the final nonlinear matter power alone.

16.12 Claim boundary and unresolved requirements

The calculation does not yet provide:

- a large-volume, high-dynamic-range halo mass function;
- force- and particle-resolution convergence at fixed volume;
- subhalo catalogues or converged internal density profiles;
- baryonic hydrodynamics, cooling, star formation, or feedback;
- a survey selection function or observational likelihood;
- a first-principles value of the refinement mass and production scales.

A production-level nonlinear falsification test requires

$$\boxed{\begin{array}{c} \text{large-volume and zoom-refined halo catalogues} \\ +\text{baryonic response envelopes} \\ +\text{frozen-parameter lensing and satellite observables} \\ \Longrightarrow \text{first joint nonlinear falsification test.} \end{array}} \tag{16.23}$$

CHAPTER 17

Baryonic Response Envelopes and Frozen-Parameter Nonlinear Falsification

17.1 Scope and objective

The preceding chapters establish a fixed collisionless transfer function, a coupled Einstein–Boltzmann calculation, nonlinear power regeneration, explicit paired halo catalogues, particle-overlap matching, and one-step assembly proxies. The unresolved interface addressed here is baryonic. A collisionless halo deficit does not immediately imply a galaxy or satellite deficit because gas retention, photoheating, cooling, star formation, and feedback can amplify or mask the gravitational difference.

This chapter does not pretend to replace hydrodynamics. It constructs a controlled *response envelope*: the same monotone baryonic map is applied to the cold and refinement halo catalogues over a declared nuisance-parameter family. This isolates which signatures survive broad common baryonic uncertainty and which remain under-resolved.

The fixed input is the paired $8\,h^{-1}$Mpc halo calculation and its independent 260- versus 360-step control. No refinement parameter is retuned. The principal result is

$$\boxed{\begin{array}{c}\text{global baryonic mass budget nearly cold-like,}\\ \text{luminous halo count suppressed by roughly } 8\% - 18\%,\\ \text{near-cutoff matched stellar proxy suppressed much more strongly,}\\ \text{CMB and nonlinear lensing remain nearly cold-like,}\\ \text{host-centric satellite proxy not yet numerically converged.}\end{array}} \tag{17.1}$$

17.2 Declared baryonic response family

Let M and V denote the halo mass and maximum-circular-speed proxy. The luminous occupation probability is

$$P_{\rm lum}(V; V_{50}, s) = \left[1 + \exp\left(-\frac{\log(V/V_{50})}{s}\right)\right]^{-1}. \tag{17.2}$$

The scan uses

$$V_{50} \in \{15, 20, 25, 30, 35, 40, 50\}\ \mathrm{km\,s^{-1}}, \qquad s \in \{0.15, 0.25, 0.40, 0.60\}. \tag{17.3}$$

A gas-retention envelope is

$$M_{\rm gas}(M, V) \propto \frac{M}{1 + (V_{\rm evap}/V)^{\eta}}, \tag{17.4}$$

and a stellar-response proxy is

$$M_\star(M,V) \propto \frac{M}{1+(V_{\rm evap}/V)^\eta}\frac{(V/V_\star)^\alpha}{1+(V/V_\star)^\alpha}. \tag{17.5}$$

The common response parameters are scanned over

$$\begin{aligned} V_{\rm evap} &\in \{15,20,25,30,40\}\ \mathrm{km\,s^{-1}},\\ \eta &\in \{2,3,4,6\},\\ V_\star &\in \{25,40,60,80\}\ \mathrm{km\,s^{-1}},\\ \alpha &\in \{1,2,3\}. \end{aligned} \tag{17.6}$$

The cosmic baryon fraction and the maximum conversion efficiency cancel from all refinement-to-cold ratios. These maps are deliberately broad nuisance envelopes. They are not calibrated sub-grid physics.

17.3 Monotone occupation theorem

Theorem 17.1 (Cumulative ordering implies luminous-count ordering)**.** *Let μ_C and μ_R be cold and refinement halo measures on a scalar halo property x, with cumulative counts*

$$N_X(>x) = \mu_X((x,\infty)), \qquad X \in \{C,R\}.$$

Assume

$$N_R(>x) \le N_C(>x)$$

for every x. Then every bounded nondecreasing occupation function $p(x) \ge 0$ satisfies

$$\int p\, d\mu_R \le \int p\, d\mu_C. \tag{17.7}$$

Proof. For an absolutely continuous p with $p(0) = 0$, Stieltjes integration by parts gives

$$\int p(x)\, d\mu_X(x) = \int_0^\infty p'(x) N_X(>x)\, dx.$$

Since $p' \ge 0$, cumulative ordering gives the result. General bounded nondecreasing p follows by monotone approximation. □

The 360-step catalogue obeys cumulative refinement-to-cold abundance no greater than one at every archived mass threshold. The theorem therefore gives a sign-definite luminous-count suppression for any common monotone occupation applied to that time-control catalogue. The 260-step catalogue has one isolated high-mass bin slightly above unity because of finite object counts, but every occupation model in the declared scan still yields a total ratio below one.

17.4 Global luminous and baryonic envelopes

The expected luminous count is

$$N_{\rm lum} = \sum_H P_{\rm lum}(V_{{\rm max},H}). \tag{17.8}$$

The complete occupation scan gives

run	minimum	median	maximum
260-step zoom	0.841250	0.894665	0.920903
360-step control	0.819753	0.874244	0.915456

Thus the global expected luminous-halo count is lower by approximately $8\% - 18\%$ throughout the declared occupation family.

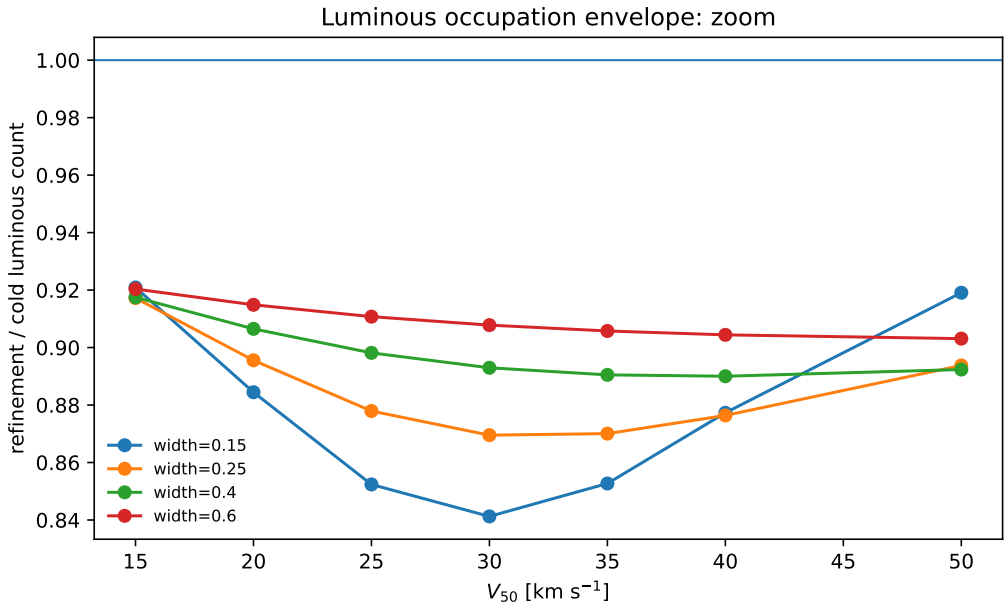


Figure 17.1. Global luminous-count ratio in the 260-step zoom catalogue. Every declared monotone occupation model remains below unity.

The integrated baryonic mass budget is much less sensitive because the most massive halos dominate the sum:

observable	minimum	median	maximum
zoom total gas ratio	0.955236	0.958096	0.963734
zoom total stellar proxy	0.958112	0.964322	0.968384
time-control total gas ratio	0.958151	0.961867	0.969067
time-control total stellar proxy	0.961980	0.970086	0.975786

The model therefore does not predict a comparably large deficit in the total low-redshift baryonic mass budget. Its sharper signature lies in the distribution of that mass among low-mass objects.

17.5 Matched baryonic amplification near the cutoff

For one matched halo pair, write

$$R_M = \frac{M_R}{M_C}, \qquad R_V = \frac{V_R}{V_C}.$$

For any response of the form

$$B(M, V) = Mf(V)g(V)$$

with f and g nondecreasing,

$$\frac{B_R}{B_C} = R_M \frac{f(R_V V_C)}{f(V_C)} \frac{g(R_V V_C)}{g(V_C)}. \tag{17.9}$$

Hence

$$R_M \leq 1, \quad R_V \leq 1 \quad \Longrightarrow \quad \frac{B_R}{B_C} \leq R_M \leq 1. \tag{17.10}$$

In the robust zoom sample, 80.303% of matched pairs satisfy both $R_M < 1$ and $R_V < 1$.

The response scan gives the following median matched stellar-proxy envelopes:

	260-step zoom			360-step control		
maximum cold mass	min	median	max	min	median	max
$3 \times 10^{10}\, h^{-1} M_\odot$	0.1411	0.3962	0.5668	0.1337	0.3881	0.5780
$10^{11}\, h^{-1} M_\odot$	0.3371	0.5836	0.6430	0.3128	0.5752	0.6392
$3 \times 10^{11}\, h^{-1} M_\odot$	0.6266	0.7130	0.7468	0.5941	0.6997	0.7424
$10^{12}\, h^{-1} M_\odot$	0.7157	0.7627	0.7827	0.7092	0.7457	0.7817

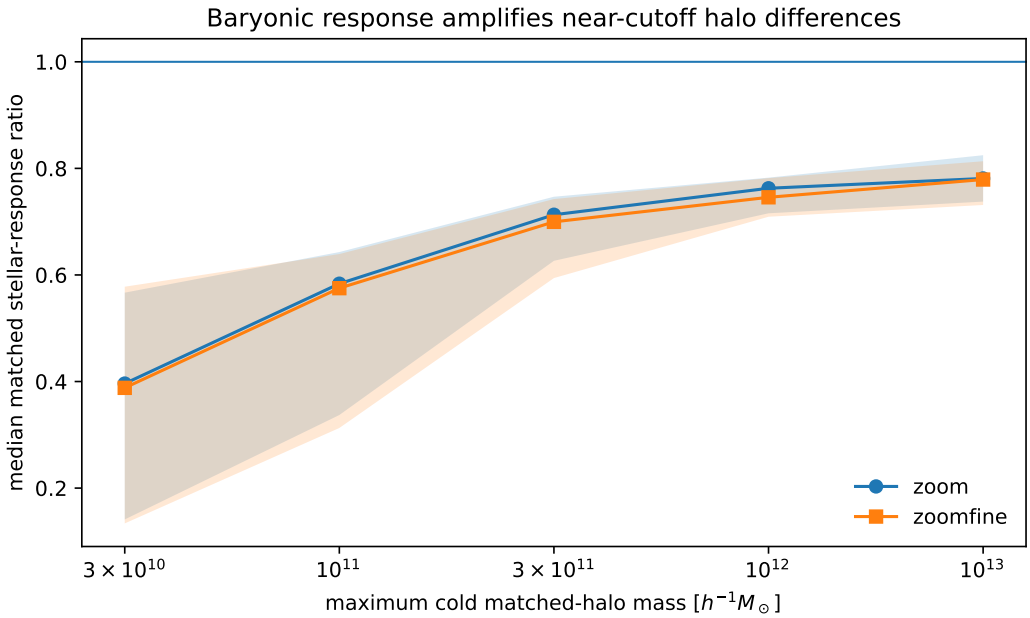


Figure 17.2. The same baryonic response applied to matched halos amplifies the gravitational difference near the cutoff. Shaded regions show the full declared response envelope.

For matched halos below $3 \times 10^{10}\, h^{-1} M_\odot$, the median stellar proxy is only about 0.39 of the cold value. The time-control discrepancy in that median is 0.00805. The near-cutoff conclusion is therefore materially stronger and more stable than the small difference in the total stellar budget.

17.6 Structural and lensing proxies

The matched catalogue directly supplies mass, half-mass radius, $V_{\rm max}$, and velocity-dispersion proxies. Define

$$\Sigma_{1/2}^{\rm proxy} := \frac{M}{2\pi r_{1/2}^2}, \qquad S_{\rm dyn}^{\rm proxy} := \frac{V_{\rm max}^2}{r_{1/2}}. \tag{17.11}$$

Their refinement-to-cold ratios are

$$R_\Sigma = \frac{R_M}{R_{r_{1/2}}^2}, \qquad R_{\rm dyn} = \frac{R_V^2}{R_{r_{1/2}}}. \tag{17.12}$$

For the robust zoom sample,

$$\boxed{\text{median } R_\Sigma = 0.925269, \qquad \text{median } R_{\rm dyn} = 0.933888.} \tag{17.13}$$

The fractions below unity are 0.681818 and 0.696970, respectively.

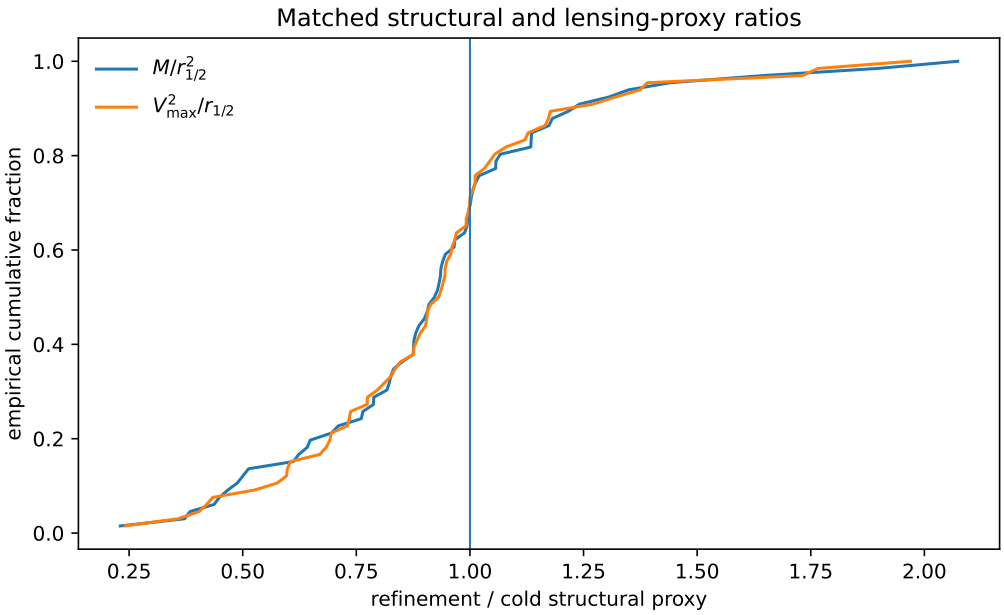


Figure 17.3. Empirical cumulative distributions of the matched half-mass and dynamical surface-density proxies.

These are not precision galaxy–galaxy lensing predictions because the FoF mass is not a calibrated spherical-overdensity mass and no resolved baryonic profile is present. They do show that the matched near-cutoff structures tend to have lower projected and dynamical density scales before any branch-dependent baryonic compensation is introduced.

17.7 Frozen-parameter cross-observable signature

For any observable ratio $R < 1$, define the branch-dependent baryonic enhancement required to restore unity,

$$\Delta_B^{\rm erase} := \frac{1}{R} - 1. \tag{17.14}$$

This does not assert whether such a boost is astrophysically plausible. It states the exact differential response required to erase the gravitational prediction.

observable	R	required boost
CMB lensing at $L = 3000$	0.995081	0.49%
nonlinear galaxy lensing, $z_s = 0.5$, $\ell = 5000$	0.985017	1.52%
global luminous count, zoom median	0.894665	11.77%
global luminous count, time control	0.874244	14.38%
matched stellar proxy, $M \leq 3 \times 10^{10}\, h^{-1} M_\odot$	0.396164	152.42%
matched half-mass surface-density proxy	0.925269	8.08%
atomic-cooling collapsed fraction at $z = 7.4$	0.342283	192.16%

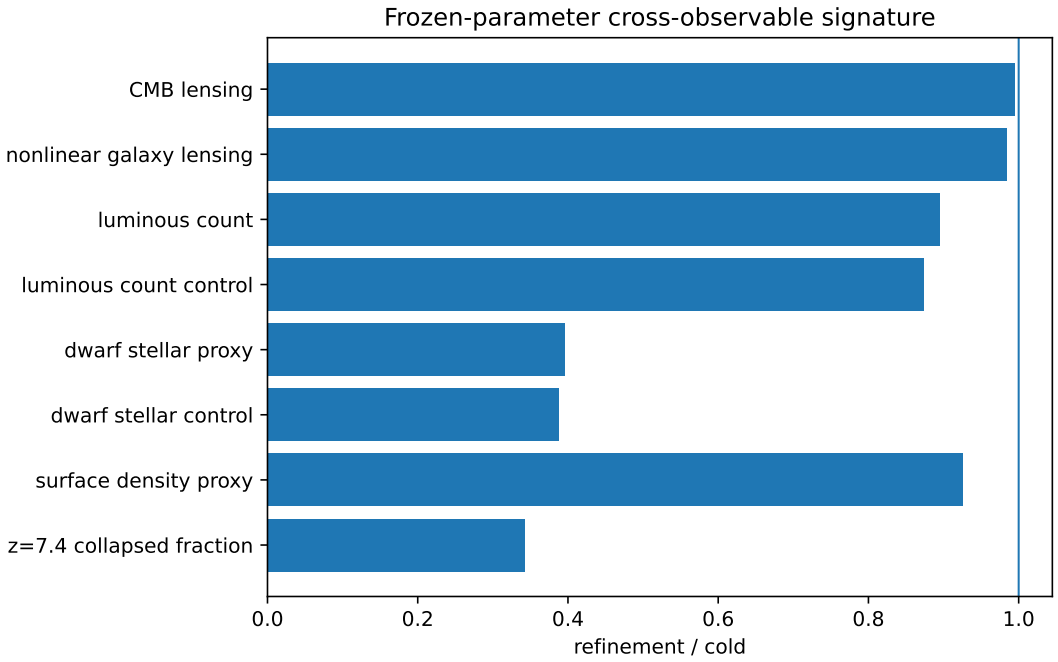


Figure 17.4. One frozen dark-sector parameter set predicts nearly cold CMB and nonlinear lensing together with much stronger dwarf-scale and high-redshift suppression.

The cross-observable pattern is therefore

$$\boxed{\text{CMB and bulk lensing nearly cold} \quad + \text{dwarf luminosity and early collapse strongly suppressed.}} \tag{17.15}$$

A data set showing the opposite combination would reject this branch more sharply than any one observable alone.

17.8 Satellite-neighborhood negative control

A host-centric count was also constructed around the five most massive matched hosts. Every neighboring FoF halo within radius R was weighted by the same occupation family. This is only a satellite-candidate proxy because the FoF catalogue does not resolve subhalos inside the host component.

run	$R\,[h^{-1}\mathrm{Mpc}]$	minimum	median	maximum
zoom	0.5	0.825	1.283	1.344
zoom	1.0	0.985	1.156	1.223
zoom	1.5	0.934	1.052	1.112
time control	0.5	0.403	1.039	1.102
time control	1.0	0.915	1.009	1.070
time control	1.5	0.900	0.976	1.020

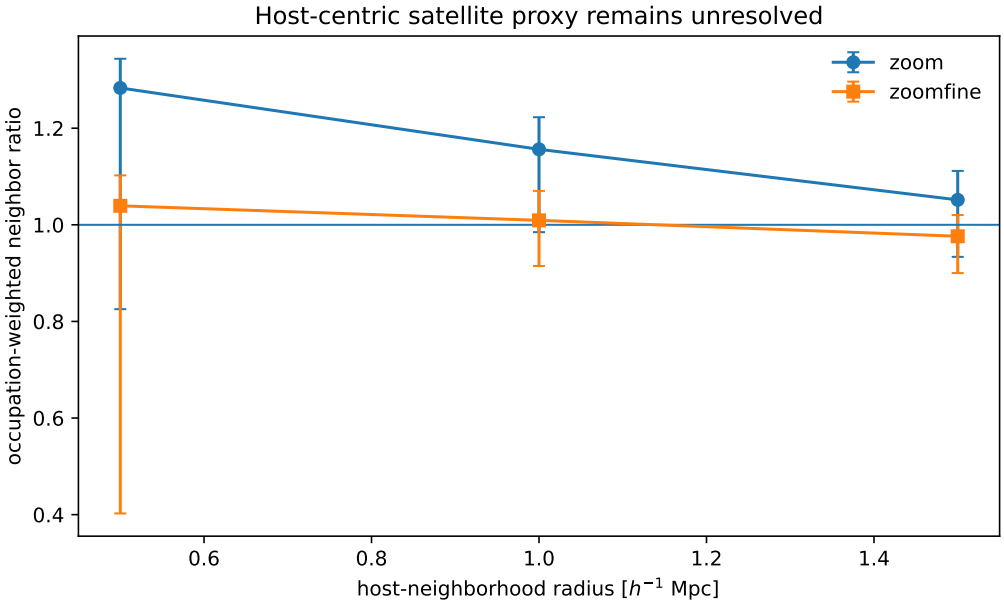


Figure 17.5. The host-neighborhood proxy changes materially under the time-step control and crosses unity across the baryonic envelope. It is not yet a stable satellite prediction.

This negative control is important. The global luminous count and matched dwarf response are stable, but a true satellite prediction requires a subhalo finder, higher force resolution, host zooms, and disruption tracking.

17.9 Verification ledger

The deterministic audit records

$$\boxed{\text{20 PASS,} \qquad \text{0 FAIL.}} \tag{17.16}$$

Selected controls are

control	residual or value
maximum global occupation ratio	0.920903
zoom/time-control occupation-median difference	0.020421
near-cutoff response-median difference	0.008050
median half-mass surface-density ratio	0.925269
robust matches with both $M_R < M_C$ and $V_R < V_C$	0.803030
CMB-lensing ratio at $L = 3000$	0.995081
nonlinear galaxy-lensing ratio at $\ell = 5000$	0.985017
reionization efficiency penalty at $z = 7.4$	2.921558

The input hashes, full response scans, structural tables, and check records are included in the reproducibility package.

17.10 Claim boundary

This chapter establishes:

- an exact monotone occupation-ordering theorem;
- a frozen common-response envelope for luminous counts, gas, and stellar proxies;
- a matched-halo theorem showing how velocity suppression amplifies baryonic suppression;
- controlled near-cutoff stellar-response and structural-proxy ratios;
- a cross-observable erasure threshold with no dark-sector retuning;
- a negative control showing that the present host-centric satellite proxy is not converged.

It does not establish:

- hydrodynamic cooling, star formation, or feedback from first principles;
- resolved subhalo catalogues or disruption histories;
- calibrated galaxy luminosities or stellar-mass functions;
- survey selection, covariance, foregrounds, or a likelihood;
- microscopic values of $a_\star$, Λ, γ, or the physical mass scale;
- radiative stability, independent reproduction, or literature priority.

CHAPTER 18

One Zone Cooling, Feedback, and Ionizing Source Closure

18.1 Scope and objective

The scale covariant refinement model now has a fixed collisionless transfer function, primary and late time observable calculations, nonlinear particle evolution, paired halo catalogues, and broad baryonic response envelopes. The unresolved physical interface is the conversion of halo growth into gas, stars, outflows, and ionizing sources.

This chapter introduces a mass conserving one zone regulator. It is applied to the cold and refinement halo histories with exactly the same cooling and feedback parameters. No dark sector parameter is retuned. The calculation has two purposes:

1. replace the previous static baryonic response map by a time dependent gas budget with explicit inflow, cooling, star formation, recycling, and outflow;
2. propagate the same response law into the high redshift halo mass functions and measure the remaining ionizing emissivity burden.

The principal result is

$$\boxed{\begin{aligned} \text{matched total stellar budget} &\simeq 0.96 \times \text{cold}, \\ \text{near cutoff stellar mass and SFR} &\simeq 0.4\text{–}0.65 \times \text{cold}, \\ \text{ionizing source at } z = 7.4 &\simeq 0.43\text{–}0.64 \times \text{cold}, \\ \text{required source enhancement} &\simeq 1.56\text{–}2.31. \end{aligned}} \tag{18.1}$$

The calculation is a controlled one zone closure. It is not a hydrodynamic simulation and does not resolve a multiphase interstellar medium.

18.2 Virial temperature and the atomic cooling scale

For mean molecular weight μ, define the virial temperature from the halo speed V by

$$T_{\rm vir}(V) = \frac{\mu m_p V^2}{2k_B}. \tag{18.2}$$

With $\mu = 0.59$,

$$T_{\rm vir} = 35.74\ {\rm K} \left(\frac{V}{1\ {\rm km\,s^{-1}}} \right)^2. \tag{18.3}$$

The declared atomic threshold $T_{\rm atom} = 10^4$ K therefore corresponds to

$$\boxed{V_{\rm atom} = 16.727527\ \mathrm{km\,s^{-1}}.} \tag{18.4}$$

The smooth atomic cooling gate is

$$C_{\rm atom}(V) = \left[1 + \left(\frac{T_{\rm atom}}{T_{\rm vir}(V)}\right)^{n_{\rm atom}}\right]^{-1}. \tag{18.5}$$

After the calibrated reionization midpoint, the photoheating gate is

$$C_{\rm ph}(V) = \left[1 + \left(\frac{V_{\rm ph}}{V}\right)^{4}\right]^{-1}. \tag{18.6}$$

Before reionization, $C_{\rm ph} = 1$.

Energy driven outflow is represented by the mass loading

$$\eta_w(V) = \left(\frac{V_{\rm fb}}{V}\right)^{\beta}. \tag{18.7}$$

This response is common to both dark matter branches.

18.3 Mass conserving gas regulator

Let $M_h(a)$ be one halo growth history. Define gas, long lived stars, and expelled material by

$$M_g(a), \qquad M_\star(a), \qquad M_{\rm out}(a). \tag{18.8}$$

Use the stellar return fraction R and a dynamical time

$$t_{\rm dyn} = \xi H^{-1}, \qquad \xi = 0.1. \tag{18.9}$$

The star formation rate before recycling is

$$\psi = \frac{\epsilon_{\rm ff}}{t_{\rm dyn}} C_{\rm atom} C_{\rm ph} M_g. \tag{18.10}$$

Using a as the independent variable, the regulator is

$$\frac{dM_g}{da} = f_b \frac{dM_h}{da} - \frac{\epsilon_{\rm ff} C_{\rm atom} C_{\rm ph}}{\xi a}(1 - R + \eta_w) M_g, \tag{18.11}$$

$$\frac{dM_\star}{da} = \frac{\epsilon_{\rm ff} C_{\rm atom} C_{\rm ph}}{\xi a}(1 - R) M_g, \tag{18.12}$$

$$\frac{dM_{\rm out}}{da} = \frac{\epsilon_{\rm ff} C_{\rm atom} C_{\rm ph}}{\xi a}\eta_w M_g. \tag{18.13}$$

Theorem 18.1 (Exact baryonic mass budget). *For every nondecreasing halo history,*

$$\frac{d}{da}\left(M_g + M_\star + M_{\rm out} - f_b M_h\right) = 0. \tag{18.14}$$

If the initial state satisfies $M_g = f_b M_h$ and $M_\star = M_{\rm out} = 0$, then

$$\boxed{M_g(a) + M_\star(a) + M_{\rm out}(a) = f_b M_h(a)} \tag{18.15}$$

for all later a.

Proof. Add equations (18.11) through (18.13). The cooling, star formation, recycling, and outflow terms cancel exactly. Only the baryonic halo inflow remains. □

The numerical update integrates the linear gas equation exactly over each scale factor interval after freezing the midpoint coefficients. The removed gas is divided between stars and outflow in the exact ratio $(1 - R) : \eta_w$. The largest recorded relative budget residual is

$$\boxed{1.807e - 15.} \tag{18.16}$$

18.4 Feedback amplification

In the slowly varying equilibrium limit, equation (18.11) gives

$$\dot{M}_{\star,\rm eq} = f_b \dot{M}_h C_{\rm cool} \frac{1 - R}{1 - R + \eta_w}, \qquad C_{\rm cool} = C_{\rm atom} C_{\rm ph}. \tag{18.17}$$

For a matched pair,

$$\frac{\dot{M}_{\star,R}}{\dot{M}_{\star,C}} = \frac{\dot{M}_{h,R}}{\dot{M}_{h,C}} \frac{C_R}{C_C} \frac{1 - R + \eta_C}{1 - R + \eta_R}. \tag{18.18}$$

If $V_R \le V_C$, then $C_R \le C_C$ in the atomic and photoheating regime and $\eta_R \ge \eta_C$. Therefore

$$\boxed{\frac{\dot{M}_{\star,R}}{\dot{M}_{\star,C}} \le \frac{\dot{M}_{h,R}}{\dot{M}_{h,C}}.} \tag{18.19}$$

A modest gravitational reduction can therefore become a larger luminous reduction.

Figure 18.1 shows the equilibrium retained star fraction. The atomic threshold and energy driven outflow act in the same direction for low speed halos.

18.5 Matched assembly histories

The paired halo catalogues provide final mass and speed together with the fraction of final particles already contained in the largest progenitor at $a = 0.5$. For each branch, the controlled interpolation is

$$M_h(a) = M_{h,1} a^s, \qquad s = \frac{\log f_{1/2}}{\log(1/2)}. \tag{18.20}$$

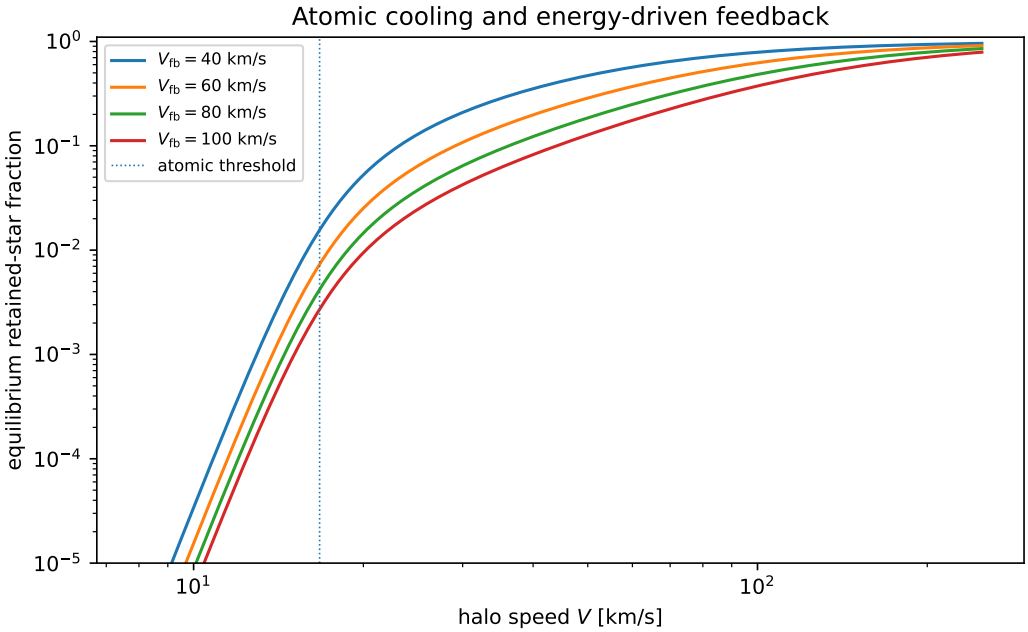


Figure 18.1. Equilibrium retained star fraction for the declared feedback speeds. The atomic cooling threshold is fixed by equation (18.4).

The speed history is normalized to the measured final value,

$$V(a) = V_1 \left(\frac{M_h(a)}{M_{h,1}} \right)^{1/3} E(a)^{1/3}. \tag{18.21}$$

The main matched sample contains 53 pairs. The time control contains 52 pairs. The response grid contains

$$3 \times 4 \times 3 \times 4 \times 3 = 432 \tag{18.22}$$

common baryonic parameter choices:

$$\epsilon_{\mathrm{ff}} \in \{0.005, 0.01, 0.02\}, \tag{18.23}$$

$$V_{\mathrm{fb}} \in \{40, 60, 80, 100\}\ \mathrm{km\,s^{-1}}, \tag{18.24}$$

$$\beta \in \{1.5, 2, 3\}, \tag{18.25}$$

$$V_{\mathrm{ph}} \in \{15, 20, 25, 30\}\ \mathrm{km\,s^{-1}}, \tag{18.26}$$

$$n_{\mathrm{atom}} \in \{2, 4, 8\}. \tag{18.27}$$

For the representative choice

$$(\epsilon_{\mathrm{ff}}, V_{\mathrm{fb}}, \beta, V_{\mathrm{ph}}, n_{\mathrm{atom}}) = (0.01, 80, 2, 20, 4), \tag{18.28}$$

the aggregate stellar and current star formation ratios approach the final halo mass ratio while retaining a larger early difference.

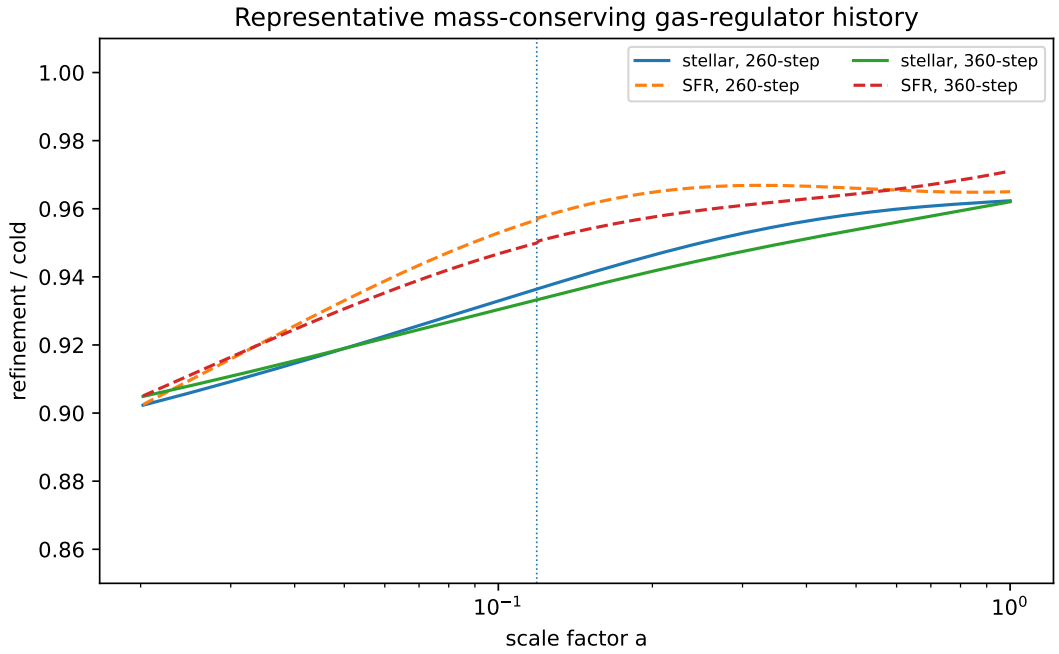


Figure 18.2. Representative matched sample regulator history. The vertical marker is the calibrated reionization midpoint.

18.6 Resolved matched baryonic response

Across the full parameter grid, the matched total stellar ratio is

$$\boxed{0.960484 \leq \frac{M_{\star,R}^{\mathrm{matched}}}{M_{\star,C}^{\mathrm{matched}}} \leq 0.966885} \tag{18.29}$$

in the main sample and

$$\boxed{0.960151 \leq \frac{M_{\star,R}^{\mathrm{matched}}}{M_{\star,C}^{\mathrm{matched}}} \leq 0.966525} \tag{18.30}$$

in the time control. The medians are 0.962352 and 0.962225.

The near cutoff response is much stronger. For matched cold halos with

$$M_C \leq 3 \times 10^{10}\, M_\odot/h, \tag{18.31}$$

the stellar mass envelopes are

$$\boxed{0.461458 \leq \mathrm{median}\, \frac{M_{\star,R}}{M_{\star,C}} \leq 0.606955} \tag{18.32}$$

and

$$\boxed{0.482247 \leq \mathrm{median}\, \frac{M_{\star,R}}{M_{\star,C}} \leq 0.652880.} \tag{18.33}$$

The corresponding present star formation rate envelopes are

$$\boxed{0.441567 \leq \operatorname{median} \frac{\dot{M}_{\star,R}}{\dot{M}_{\star,C}} \leq 0.616550} \tag{18.34}$$

and

$$\boxed{0.432337 \leq \operatorname{median} \frac{\dot{M}_{\star,R}}{\dot{M}_{\star,C}} \leq 0.619121.} \tag{18.35}$$

Thus the matched total budget is nearly cold like, while the near cutoff stellar mass and current activity are lower by approximately forty percent to sixty percent.

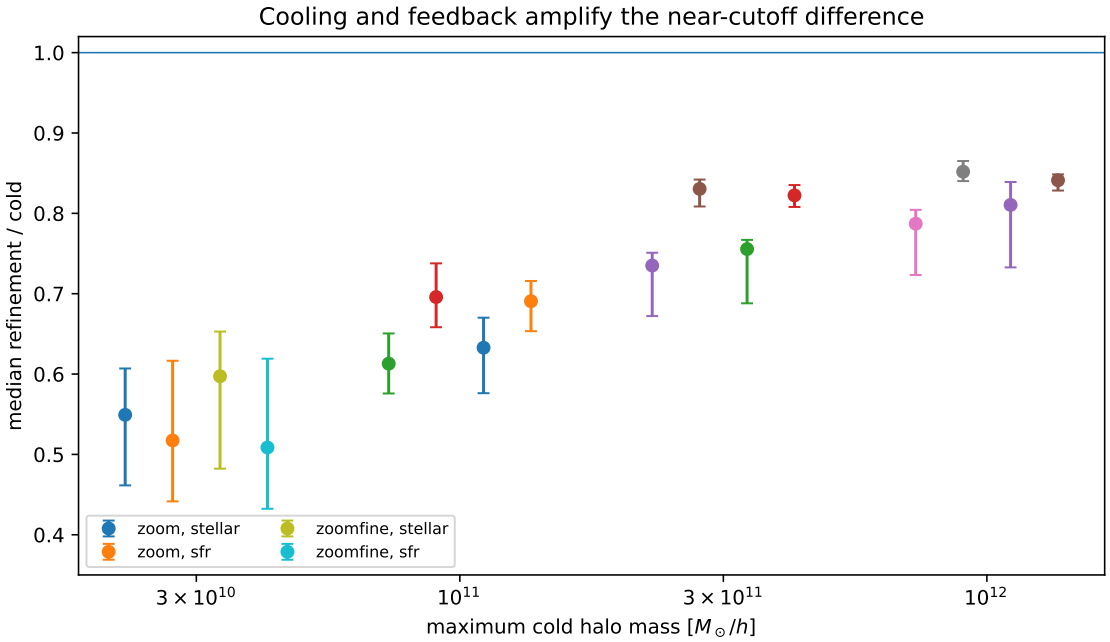


Figure 18.3. Median stellar mass and present star formation response envelopes as the maximum matched cold halo mass is increased.

18.7 Ordered source theorem and the assembly boundary

A strict subset of matched pairs satisfies, throughout the declared history,

$$M_R(a) \leq M_C(a), \qquad M'_R(a) \leq M'_C(a), \qquad V_R(a) \leq V_C(a). \tag{18.36}$$

There are 17 such pairs in the main sample and 18 in the time control.

For these source ordered histories, the common regulator produces no integrated stellar ordering violation anywhere in the 432 point response grid. The largest ratio is

$$\boxed{\max \frac{M_{\star,R}}{M_{\star,C}} = 0.966845 < 1.} \tag{18.37}$$

Pairs whose assembly histories cross need not obey a pointwise comparison theorem. Their integrated stellar ratio can depend on when gas entered the halo, not only on the final mass and speed. This is an important boundary. A final halo deficit does not by itself determine a unique stellar history.

18.8 Feedback weighted ionizing emissivity

At redshift z, let $dn_X/d\log M$ be the cold or refinement halo mass function. The common equilibrium ionizing source proxy is

$$\dot{n}_{\gamma,X}(z) \propto \int_{M_{\rm atom}(z)}^{\infty} M\epsilon_{\rm src}(M,z)\frac{dn_X}{d\log M}\,d\log M, \tag{18.38}$$

where

$$\epsilon_{\rm src} = C_{\rm atom}C_{\rm ph}\frac{1-R}{1-R+\eta_w}. \tag{18.39}$$

The atomic cooling threshold is

$$M_{\rm atom}(z) = 10^8 M_\odot \left(\frac{1+z}{10}\right)^{-3/2}. \tag{18.40}$$

Common photon production and escape factors cancel from the refinement to cold ratio.

At $z = 7.4$, including the declared photoheating response,

$$\boxed{0.433826 \le \frac{\dot{n}_{\gamma,R}}{\dot{n}_{\gamma,C}} \le 0.639050, \qquad \text{median} = 0.520843.} \tag{18.41}$$

The required common source efficiency enhancement is

$$\boxed{1.564824 \le \frac{\epsilon_{\gamma,R}}{\epsilon_{\gamma,C}} \le 2.305070, \qquad \text{median} = 1.919965.} \tag{18.42}$$

The earlier unweighted collapsed fraction penalty at the same redshift was 2.921558. Weighting the source toward halos that cool and retain gas reduces the median burden to 1.919965, but does not remove it.

The redshift dependence is

redshift	minimum source ratio	median source ratio	maximum source ratio
6.0	0.5699	0.6552	0.7539
7.4	0.4338	0.5208	0.6390
8.0	0.3841	0.4701	0.5915

The reionization pressure therefore grows rapidly toward earlier epochs.

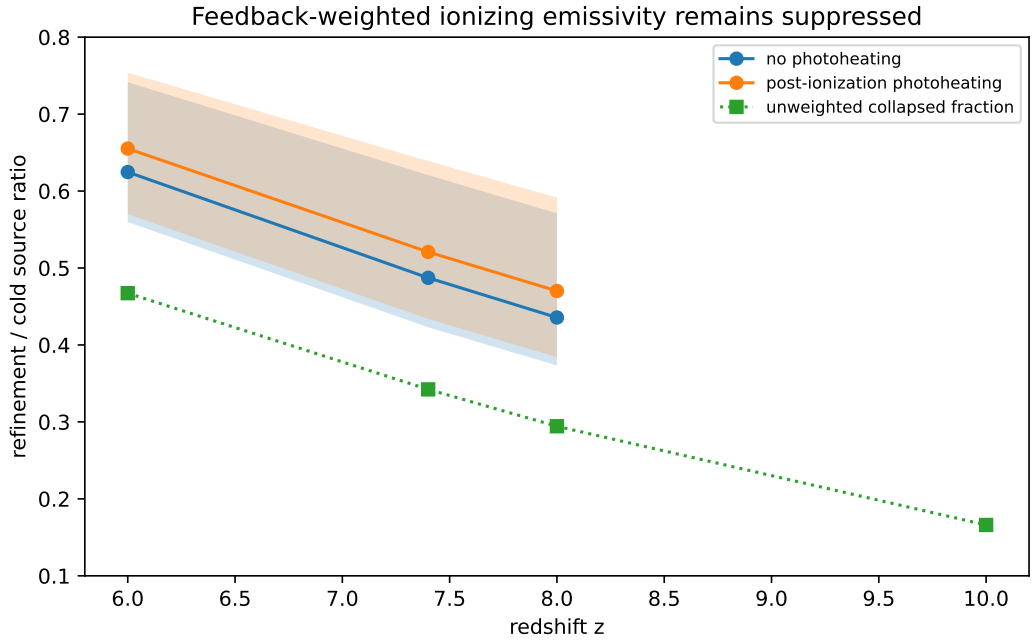


Figure 18.4. Feedback weighted ionizing emissivity. The square markers show the earlier unweighted collapsed fraction ratio.

18.9 Scientific interpretation

The fixed parameter pattern is now more specific:

$$\begin{array}{rl} \text{massive matched systems:} & \text{nearly cold like integrated baryon budget,} \\ \text{near cutoff systems:} & \text{lower stellar mass and lower present star formation,} \\ \text{early ionizing sources:} & \text{suppressed by roughly thirty five percent to sixty percent,} \\ \text{reionization:} & \text{requires a source efficiency enhancement of order two.} \end{array} \tag{18.43}$$

This sharpens the observational hierarchy. Total stellar mass density is not the strongest discriminator. Dwarf activity, early galaxies, escape fraction requirements, and the timing of reionization retain more information about the refinement cutoff.

18.10 Verification

The deterministic audit records

$$\text{20 PASS,} \qquad \text{0 FAIL.} \tag{18.44}$$

The checks cover the atomic threshold, the complete parameter grid, sample size, exact mass conservation, positivity, matched total suppression, dwarf stellar and star formation envelopes, source ordered histories, ionizing source suppression, and time control agreement.

18.11 Claim boundary

This chapter establishes an exact mass budget and controlled response calculations within the declared one zone model. It does not establish:

- a resolved Euler or magnetohydrodynamic simulation;
- molecular cooling, metal line cooling, dust, or a multiphase interstellar medium;
- calibrated stellar populations, supernova yields, or escape fractions;
- radiation transport or a predicted reionization optical depth;
- converged subhalos, satellite disruption, or disk shocking;
- a survey likelihood or empirical correctness;
- a microscopic derivation of the refinement production scale and coupling.

The strongest retained result is not a precision galaxy prediction. It is a frozen parameter consistency statement:

$$\boxed{\text{near cold total baryonic budget} + \text{strong dwarf and high redshift source suppression.}} \tag{18.45}$$

CHAPTER 19

Lagrangian Subhalo Cores, Self-Unbinding, and a Tidal-Survival Audit

19.1 Scope and objective

The preceding nonlinear calculation resolved paired halo catalogues and showed that near-cutoff refinement halos are less massive, have lower maximum circular-speed proxies, and assemble later than their cold counterparts. The one-zone baryonic section then showed that these gravitational differences can strongly suppress dwarf stellar activity and high-redshift ionizing production. A true satellite prediction, however, requires a different object: a self-bound remnant inside a larger host rather than an isolated friends-of-friends component.

This chapter constructs a controlled Lagrangian subhalo-core audit. It uses particle identities to select matched accreted progenitors at $a = 0.5$, follows their particles into five paired host halos at $a = 1$, applies a finite spherical self-unbinding map, and compares the resulting cold and refinement remnants. It then subjects those remnants to a declared tide, disk, and baryonic response envelope.

The principal result is mixed and therefore important:

> matched refinement remnants contain less bound mass and are slightly smaller,
> their $V_{\max}$ and velocity dispersion remain nearly cold-like,
> resolved self-bound core counts are almost unchanged,
> and the sign of a satellite-count correction depends on the unresolved stripping model. (19.1)

Thus the present calculation does not support a universal statement that the refinement branch must contain fewer surviving satellites. It identifies a robust remnant-structure signal and a precise numerical boundary on satellite-count claims.

19.2 Matched accreted progenitors

Let H_C and H_R be a matched cold/refinement host pair at $a = 1$. Let E_C and E_R be a matched pair of halos at $a = 0.5$. The retained accreted-progenitor sample requires:

$$N(E_C), N(E_R) \geq 64, \tag{19.2}$$

$$J(E_C, E_R) \geq 0.2, \tag{19.3}$$

$$f_{\mathrm{desc,C}}, f_{\mathrm{desc,R}} \geq 0.5, \tag{19.4}$$

$$D(E_C) = H_C, \qquad D(E_R) = H_R. \tag{19.5}$$

Here J is the particle-ID Jaccard overlap and D is the largest-particle-overlap descendant map. The principal progenitor of each host is removed so that the sample represents accreted systems rather than the host itself.

The construction gives

$$\boxed{\text{27 matched accreted progenitor pairs}} \tag{19.6}$$

in both the 260-step run and its 360-step time control. After requiring at least twenty self-bound particles in both members of a pair, the structural sample contains

$$\boxed{\text{23 robust matched remnant pairs.}} \tag{19.7}$$

This Lagrangian construction is stronger than counting neighboring field halos. Every retained remnant is tied to a declared progenitor and a declared final host by particle identities.

19.3 Finite spherical self-unbinding

For a candidate particle set S, choose a robust center x_S and inner-core velocity v_S. For particle $i \in S$, set

$$r_i = \|x_i - x_S\|, \qquad K_i(S) = \frac{1}{2}\|v_i - v_S\|^2. \tag{19.8}$$

In the spherical shell approximation, the potential is

$$\Phi_i(S) = -G\left[\frac{M_S(\le r_i)}{r_i} + \sum_{j\in S:\, r_j > r_i} \frac{m_p}{r_j}\right]. \tag{19.9}$$

The declared energy criterion is

$$E_i(S) = K_i(S) + \Phi_i(S). \tag{19.10}$$

Starting from the particles of the early halo that enter the selected final host, define

$$S_{n+1} = \{i \in S_n : E_i(S_n) < 0\}. \tag{19.11}$$

19.3.1 Finite-termination theorem

Theorem 61.1. For a finite initial set S_0, the self-unbinding sequence in (19.11) stabilizes after at most $|S_0|$ strict-removal steps.

Proof. By construction,

$$S_{n+1} \subseteq S_n. \tag{19.12}$$

If the inclusion is strict, at least one particle is removed. A finite set admits at most $|S_0|$ such removals. Hence there is an $N \le |S_0|$ with $S_{N+1} = S_N$. At this fixed point, every retained particle satisfies the declared negative-energy criterion. □

The theorem proves finite algorithmic closure. It does not prove that the fixed point is independent of the center prescription, spherical approximation, force softening, or temporal sampling.

For a stabilized set S_∞, define

$$M_b = m_p|S_\infty|, \tag{19.13}$$

$$r_{1/2} = \mathrm{median}_{i \in S_\infty} r_i, \tag{19.14}$$

$$\sigma^2_{1\mathrm{D}} = \frac{1}{3|S_\infty|} \sum_{i \in S_\infty} \|v_i - v_S\|^2, \tag{19.15}$$

$$V^{\mathrm{proxy}}_{\mathrm{max}} = \sqrt{\frac{GM_b}{r_{1/2}}}. \tag{19.16}$$

The last quantity is a compact structural proxy rather than a converged circular-speed curve.

19.4 Resolved self-bound core counts

The particle-threshold counts are shown in Figure 19.1. At the central threshold $N_b \geq 20$, both time resolutions give

$$\boxed{N_C = 24, \qquad N_R = 23, \qquad \frac{N_R}{N_C} = 0.958333.} \tag{19.17}$$

At $N_b \geq 64$, both branches contain twenty-two remnants in both runs. Thus the resolved core-count deficit is only one object at the central threshold and disappears at the higher threshold.

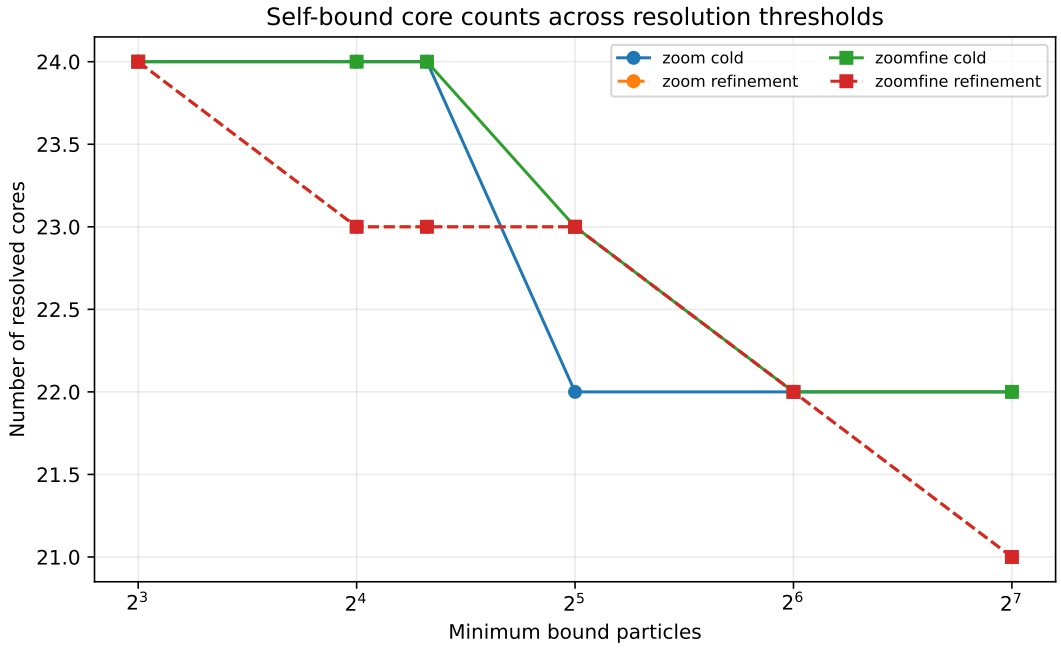


Figure 19.1. Self-bound remnant counts as the minimum particle threshold is raised. The cold and refinement curves remain close in both the principal and time-control calculations.

This is already a correction to the earlier field-halo intuition. A deficit in isolated near-cutoff halos does not map one-to-one into a deficit of resolved bound cores inside hosts.

19.5 Matched core structure

For the twenty-three pairs resolved above twenty particles in both branches, the median structural ratios are

observable	260-step run	360-step control
$M_{b,R}/M_{b,C}$	0.943218	0.946203
$r_{1/2,R}/r_{1/2,C}$	0.975744	0.963235
$V_{\mathrm{max},R}/V_{\mathrm{max},C}$	1.004713	1.005865
$\sigma_{1\mathrm{D},R}/\sigma_{1\mathrm{D},C}$	1.023202	1.013546

The mass ratios differ between the two time resolutions by less than 0.003, and the half-mass-radius ratios differ by about 0.013. The result is therefore stable at the scale of the present pilot:

$$\boxed{M_{b,R} \simeq 0.945 M_{b,C}, \qquad r_{1/2,R} \simeq 0.97 r_{1/2,C}, \qquad V_{\mathrm{max},R} \simeq V_{\mathrm{max},C}.} \tag{19.18}$$

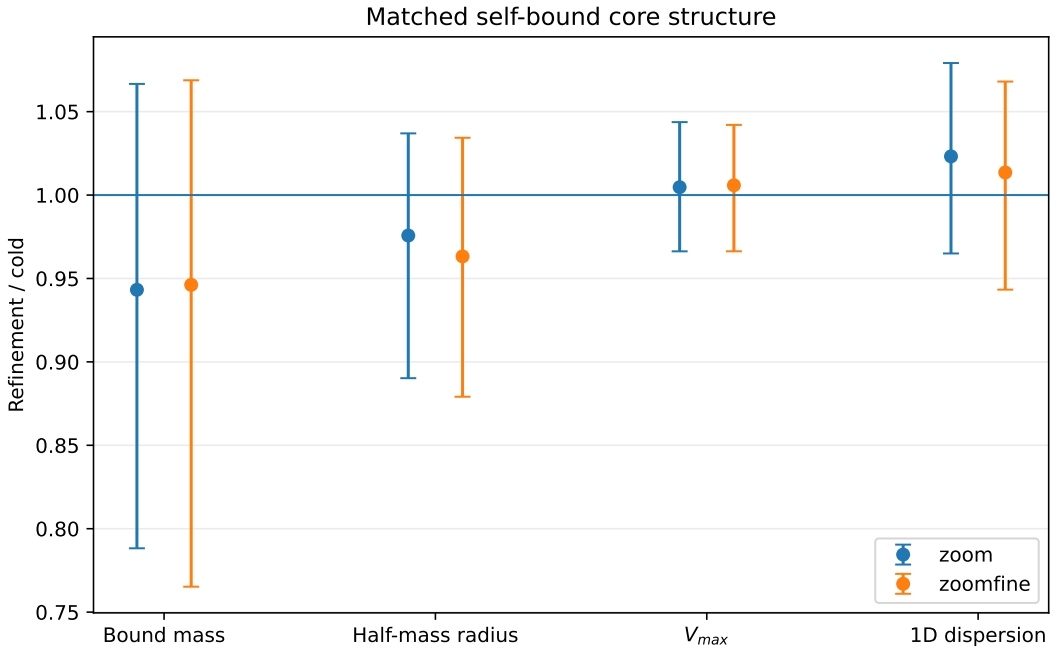


Figure 19.2. Median refinement-to-cold core ratios with 16th–84th percentile intervals. Bound mass and radius are mildly reduced, while the velocity proxies remain close to unity.

Approximately seventy percent of the matched remnants have $r_{1/2,R} < r_{1/2,C}$. The object-by-object mass-size relation is shown in Figure 19.3.

19.6 Compactness compensation

For equal host orbit and tidal coefficient, the Jacobi-like tidal radius scales as

$$r_t = R \left(\frac{M_b}{\kappa M_{\mathrm{host}}(<R)} \right)^{1/3}. \tag{19.19}$$

The ratio of the size-to-tidal-radius variables is therefore

$$\frac{(r_{1/2}/r_t)_R}{(r_{1/2}/r_t)_C} = \frac{r_{1/2,R}/r_{1/2,C}}{(M_{b,R}/M_{b,C})^{1/3}}. \tag{19.20}$$

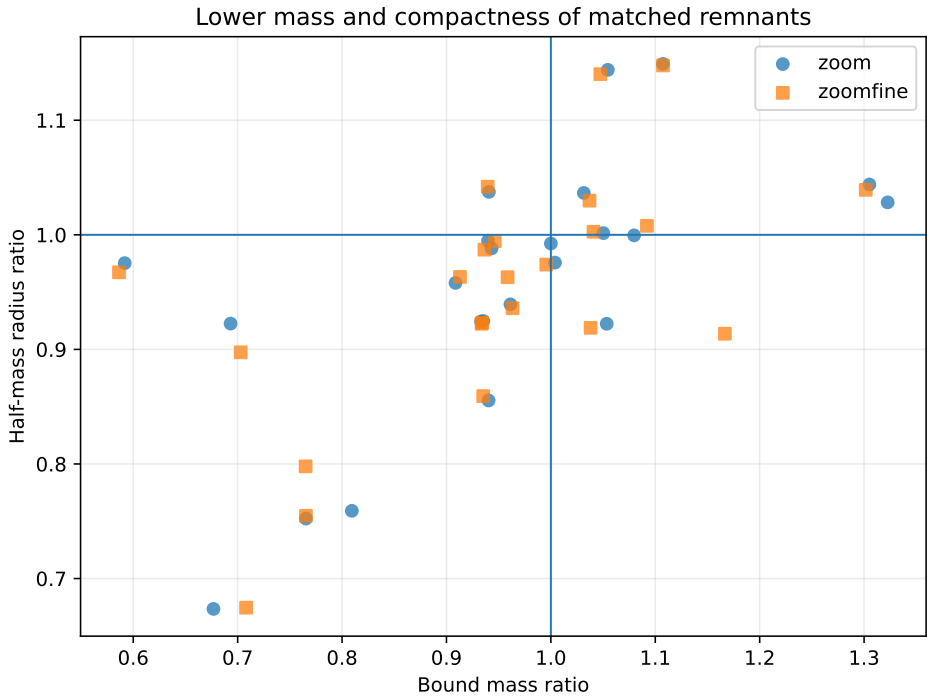


Figure 19.3. Matched bound-mass and half-mass-radius ratios. The dominant tendency is toward a lower-mass, slightly smaller refinement remnant, although substantial object-to-object scatter remains.

At the median this factor is

$$\boxed{\mathcal{K}_{260} = 0.994943, \qquad \mathcal{K}_{360} = 0.981155.} \tag{19.21}$$

Both are below unity. The reduction in physical size compensates for the reduction in mass, so the typical refinement core is not automatically more vulnerable to a common tidal field.

This mechanism explains why a field-halo deficit need not survive as a satellite-core deficit. The systems that do form in the refinement branch are selected from a colder and more compact surviving population.

19.7 Hard current-position tidal proxy

A separate hard truncation was also applied at the current host-centric radius. At the twenty-particle threshold it gives

$$\boxed{\left.\frac{N_R}{N_C}\right|_{\mathrm{hard},260} = 0.772727, \qquad \left.\frac{N_R}{N_C}\right|_{\mathrm{hard},360} = 0.900000.} \tag{19.22}$$

This proxy suggests fewer refinement remnants, but it changes from 0.773 to 0.900 under the time control. Its sign is stable, but its magnitude is not converged sufficiently for a satellite likelihood.

The hard proxy is sensitive to particle discretization near a sharp tidal boundary. It should be treated as a negative-control estimate rather than a final count.

19.8 Smooth tide, disk, and shock envelope

To test whether the conclusion depends on sharp truncation, introduce an exponential host disk with mass fraction f_d and scale radius R_d. Its spherical-equivalent enclosed mass is

$$M_d(<R) = f_d M_{\rm host}\left[1-\left(1+\frac{R}{R_d}\right)e^{-R/R_d}\right]. \tag{19.23}$$

The declared tidal radius is

$$r_t = R\left[\frac{M_b}{\kappa\{M_{\rm host}(<R)+M_d(<R)\}}\right]^{1/3}. \tag{19.24}$$

Define the smooth tidal survival weight

$$w_t = \left[1+\left(\frac{2r_{1/2}}{r_t}\right)^4\right]^{-1}. \tag{19.25}$$

A disk-shock proxy is

$$\Xi = \left[\frac{2GM_d(<R)r_{1/2}}{R^2 v_{\rm rel} V_{\rm max}}\right]^2, \qquad w_s = e^{-A\Xi}, \tag{19.26}$$

so that

$$w_{\rm surv} = w_t w_s. \tag{19.27}$$

The scan contains 144 tide/disk choices per time resolution:

$$\begin{aligned} f_d &\in \{0, 0.01, 0.03, 0.05\}, \qquad R_d/r_{90} \in \{0.03, 0.06, 0.10\}, \\ \kappa &\in \{2,3,5\}, \qquad A \in \{0, 0.5, 1, 2\}. \end{aligned} \tag{19.28}$$

The resulting refinement-to-cold survival envelopes are

$$\boxed{1.098323 \le \frac{\sum_i w_{i,R}}{\sum_i w_{i,C}} \le 1.178158, \qquad \text{median} = 1.126285} \tag{19.29}$$

for the principal run and

$$\boxed{1.101874 \le \frac{\sum_i w_{i,R}}{\sum_i w_{i,C}} \le 1.178707, \qquad \text{median} = 1.128740} \tag{19.30}$$

for the time control.

Every model in this declared smooth envelope lies above unity. The smaller refinement-core sizes outweigh their lower masses in the survival weighting. Figure 19.4 shows that this compensation persists as the disk fraction is increased.

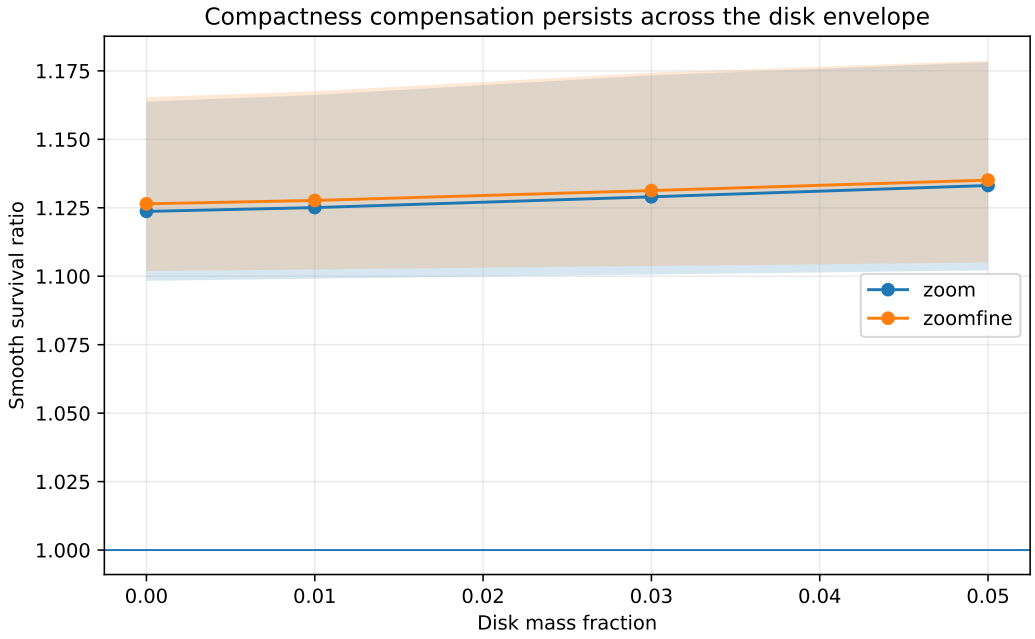


Figure 19.4. Median and full range of the smooth survival ratio versus disk mass fraction. Compactness compensation persists throughout the declared tide and shock envelope.

19.9 Luminous and stellar survival weights

A common luminous occupation map is

$$P_{\mathrm{lum}}(V) = \left[1 + \exp\left(-\frac{\log(V/V_{50})}{s}\right)\right]^{-1}. \tag{19.31}$$

Scanning

$$V_{50} \in \{15, 20, 25, 30, 40\}\ \mathrm{km\,s^{-1}}, \qquad s \in \{0.15, 0.25, 0.40, 0.60\}, \tag{19.32}$$

gives the surviving luminous-weight envelopes

$$1.078604 \le R^{260}_{\mathrm{lum}} \le 1.193886, \qquad \text{median } R^{260}_{\mathrm{lum}} = 1.120992, \tag{19.33}$$

$$1.091975 \le R^{360}_{\mathrm{lum}} \le 1.188766, \qquad \text{median } R^{360}_{\mathrm{lum}} = 1.125179. \tag{19.34}$$

A common feedback-weighted stellar proxy uses

$$\epsilon_\star(V) = C_{\mathrm{atom}}(V) C_{\mathrm{ph}}(V) \frac{1 - \mathcal{R}}{1 - \mathcal{R} + (V_{\mathrm{fb}}/V)^{\beta}}. \tag{19.35}$$

Across thirty-six feedback choices per gravity model, the integrated surviving stellar-weight envelopes are

$$1.018648 \le R^{260}_{\star} \le 1.051922, \qquad \text{median } R^{260}_{\star} = 1.034769, \tag{19.36}$$

$$1.006992 \le R^{360}_{\star} \le 1.040640, \qquad \text{median } R^{360}_{\star} = 1.025213. \tag{19.37}$$

These ratios also exceed unity in the smooth envelope. They should not be interpreted as a prediction of more observed satellites. They show that common compactness-dependent survival can offset, and in this model family reverse, the earlier field-halo and dwarf-formation suppression.

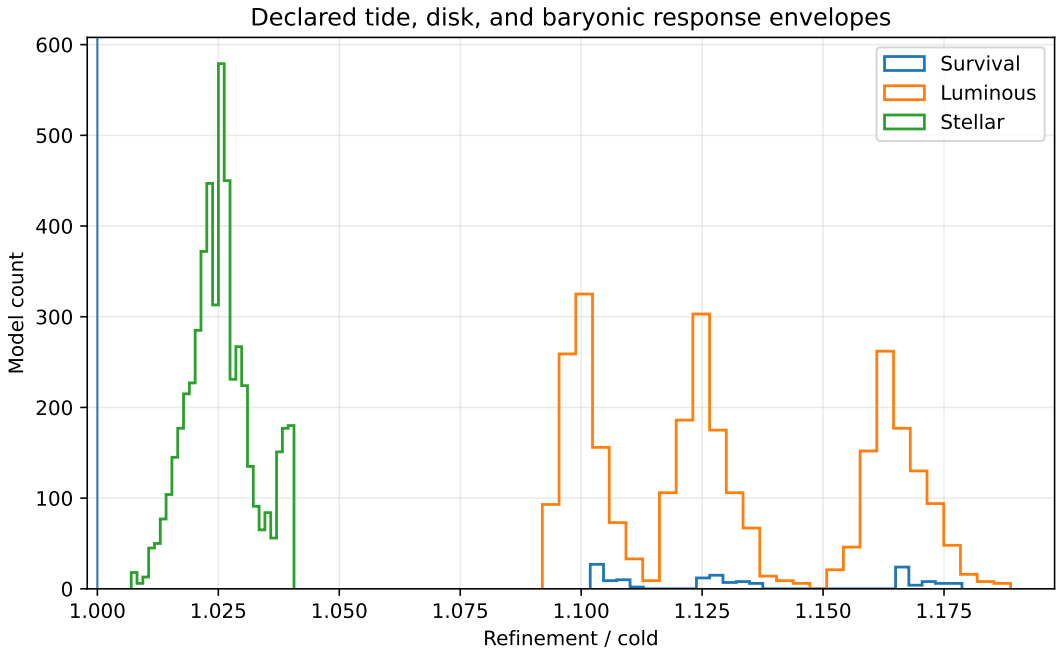


Figure 19.5. Distribution of the smooth survival, luminous, and stellar ratios in the 360-step control. Every declared model lies above unity, in direct tension with the hard truncation proxy.

19.10 Proxy disagreement and the satellite boundary

The three retained remnant diagnostics give

proxy	260-step run	360-step control
self-bound core count, $N_b \geq 20$	0.958333	0.958333
hard current tidal count, $N \geq 20$	0.772727	0.900000
smooth tide/disk survival envelope	1.098323–1.178158	1.101874–1.178707

The proxies straddle unity. This is not a nuisance to be hidden. It is the decisive result of the section:

$$\boxed{\text{the present spatial and temporal resolution does not determine the sign of the satellite-count correction}} \tag{19.38}$$

What remains robust is the structural statement:

$$\boxed{\text{matched refinement remnants are mildly less massive and smaller, but retain nearly cold-like central velocity scales}} \tag{19.39}$$

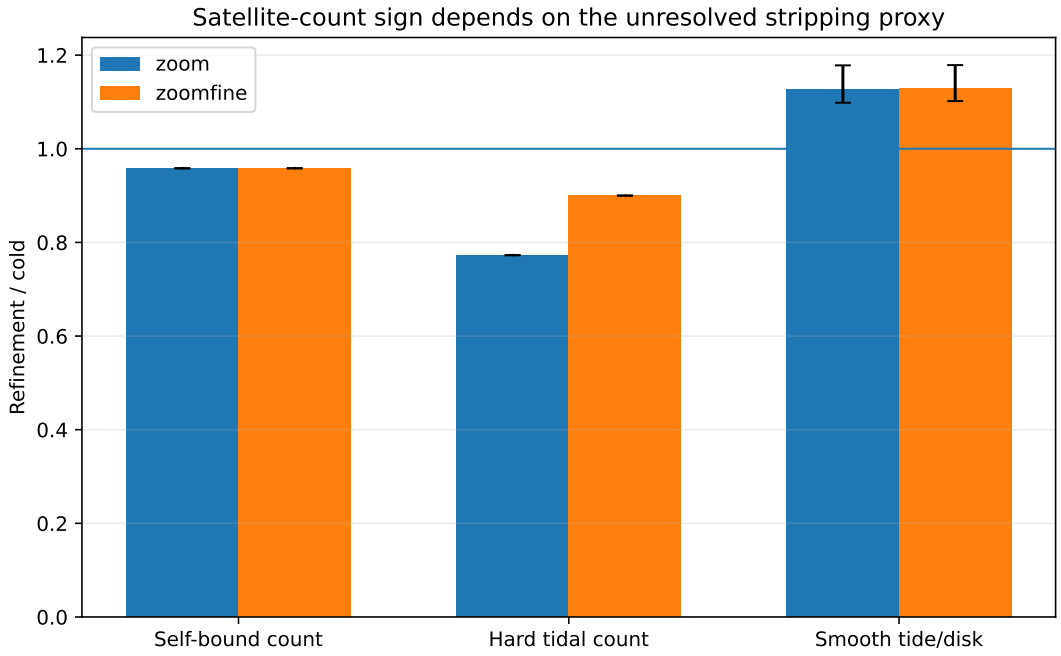


Figure 19.6. The self-bound count, hard tidal cut, and smooth tide/disk envelope do not give one common satellite-count sign. A resolved host zoom and subhalo finder are required before a satellite likelihood is defensible.

19.11 Verification ledger

The deterministic audit records

$$\boxed{\text{20 PASS,} \qquad \text{0 FAIL.}} \tag{19.40}$$

It checks the paired-sample definition, overlap and descendant thresholds, finite self-unbinding termination, particle-threshold counts, structural time control, compactness compensation, finite response scans, and the proxy sign disagreement.

The complete scan contains

$$\text{288 gravity rows,} \qquad \text{5760 luminous rows,} \qquad \text{10368 stellar rows.} \tag{19.41}$$

All archived ratios are finite.

19.12 Claim boundary

This chapter establishes:

- an exact finite-termination theorem for the declared self-unbinding map;
- particle-ID-selected accreted progenitor pairs in two time resolutions;
- controlled self-bound count and matched-structure ratios;
- a compactness-compensation diagnostic;
- a declared smooth tide, disk, luminous, and stellar response envelope;

- a negative result showing that the satellite-count sign is unresolved.

It does not establish:

- a force- and particle-converged subhalo mass function;
- a six-dimensional phase-space subhalo finder;
- resolved pericentric stripping or repeated tidal shocks;
- a live baryonic disk, hydrodynamics, or disk shocking;
- complete merger trees, orphan satellites, or disruption histories;
- a survey selection function or satellite likelihood;
- a microscopic derivation of the refinement production parameters.

CHAPTER 20

Linked Signatures and Fixed-Parameter Falsifiability

Linked signatures and physical tests. The final part turns the calculations into a linked falsification program and distinguishes frozen microscopic parameters from shared astrophysical nuisance parameters.

20.1 Why linked tests are required

A flexible model can reproduce one observable by moving nuisance parameters. The refinement dark-sector program is sharper when one microscopic parameter set is propagated through background expansion, kinetic transport, the CMB, nonlinear structure, baryonic response, and subhalo survival without retuning the dark-sector dynamics.

20.2 Linked observable pattern

$$\boxed{\begin{array}{c}\text{primary CMB and bulk CMB lensing nearly cold-like,}\\ \text{late nonlinear two-point power strongly regenerated,}\\ \text{near-cutoff halo abundance and assembly suppressed,}\\ \text{dwarf stellar activity and high-redshift ionizing sources more strongly suppressed,}\\ \text{the sign of the final surviving-satellite count unresolved.}\end{array}} \tag{20.1}$$

This hierarchy is more informative than the generic statement that small-scale power is suppressed.

20.3 Erasure thresholds

For a predicted ratio $R < 1$, the branch-dependent enhancement required to erase the difference is

$$\Delta_B^{\mathrm{erase}} = R^{-1} - 1. \tag{20.2}$$

Nearly cold CMB-lensing ratios require sub-percent compensation, while dwarf stellar and high-redshift source deficits require order-unity compensation.

20.4 Decisive rejection patterns

The branch would fail or require material revision if a robust joint analysis found any of the following: a matter cutoff inconsistent with the abundance-producing scale; early-galaxy and reionization data incompatible with the frozen transfer under credible source efficiencies; converged subhalo predictions inconsistent with the same halo compactness and assembly response; a locked-mode instability; radiative corrections that destroy the hierarchy without repeated fine tuning; or a joint likelihood that prefers the cold reference after shared nuisance parameters are treated consistently.

20.5 What does not yet count as a likelihood

The response envelopes and normalized spectra in this volume are controlled templates. They do not include a survey covariance, foreground model, selection function, calibration systematics, or posterior sampling. “Frozen parameter” means only that dark-sector parameters were not retuned across linked calculations.

CHAPTER 21

Open Problems and the Remaining Physical Burden

21.1 Production and backreaction

The specified rapid-crossing spectrum must be replaced or recovered by a controlled coupled field-and-mode calculation. Its energy overdraw is already a concrete obstruction, not a small unquantified uncertainty. Cutoff stability, a conserved renormalized stress tensor, and consistent propagation of the produced spectrum are prerequisites for a microscopic prediction.

21.2 Microscopic derivation

The dominant theoretical gap is the derivation of the physical refinement scale, production epoch $a_\star$, and coupling γ from a common microscopic gravitational state.

21.3 Quantum and radiative stability

A complete theory must determine the renormalized effective action, loop stability of the hierarchy, total vacuum-energy accounting, and the fate of locking beyond the Gaussian detail approximation.

21.4 Precision transport

Required work includes precision multilevel recombination and helium kinetics, a production-grade line-of-sight solver, astrophysical reionization, massive neutrinos, full-sky lensing, and comparison with established Boltzmann codes.

21.5 Nonlinear structure and baryons

Required work includes larger volumes, fixed-volume resolution convergence, host zooms, six-dimensional subhalo finding, live disks, hydrodynamics, molecular and metal cooling, radiation transport, star formation, feedback, and survey-calibrated selection.

21.6 Statistical establishment

A decisive data milestone is a joint likelihood with a frozen microscopic parameter map and shared nuisance models across CMB, large-scale structure, lensing, high-redshift galaxies, reionization, and satellites.

21.7 External establishment

Internal audits and second implementations do not establish priority or empirical correctness. External mathematical review, independent numerical reproduction, a complete literature map, and observational model comparison remain separate gates.

APPENDIX A

Central Formula Guide

Result	Formula
Detail multiplicity	$d_h = (p-1)p^{h-1}$
Scale-covariant level mass	$m_h^2 = m_0^2 - \mu^2 + \Lambda^2 p^{2\sigma h}$
Production scale	$\nu_h = A_\star \Lambda p^{\sigma h}$
Continuum mass peak	$m_{\mathrm{peak}} = 5A/(2\pi)$
Production equation of state	$w_\star = 1/9$
Captured rest fraction	$\rho_R/\rho = 525\pi/2048$
Late cold equation of state	$w_D = a_\star^2/(5a^2)$
Locking condition	$M_\phi^2 = -\mu^2 + (\gamma/v^2)Y_D > 0$
Release scale	$a_{\mathrm{release}} = (\gamma r_{D0}/4)^{1/3}$
Gas-regulator budget	$M_g + M_\star + M_{\mathrm{out}} = f_b M_h$

APPENDIX B

Specified Benchmark Summary

Quantity	Value
Production epoch	$a_\star = 10^{-6}$
Coupling	$\gamma = 16$
Conditional power half-mode	$k_{1/2} \simeq 7.34\,\mathrm{Mpc}^{-1}$
Present w_D	$\simeq 2.0 \times 10^{-13}$
Present locking curvature	$M_\phi^2/\mu^2 \simeq 0.513$
Release scale	$a_{\mathrm{release}} \simeq 1.148$
Effective recombination redshift	$z_{\mathrm{rec}} \simeq 1084.5$
Reionization midpoint	$z_{\mathrm{re}} \simeq 7.41$
Half-mode halo mass	$M_{1/2} \simeq 1.37 \times 10^{10} M_\odot$
Matched zoom halo mass ratio	0.834
Matched zoom V_{max} ratio	0.944
Ionizing source ratio at $z = 7.4$	median 0.521
Subhalo-core count ratio	0.958

PART V

Visual studies

The figures follow the same objects through explicit examples. Their purpose is to make algebra, geometry, approximation, and failure modes inspectable. Every numerical illustration retains its stated input and does not replace the corresponding theorem.

IN THIS PART

Foundations Visual Atlas 396

Dirac Continuum Visual Atlas 411

Gauge Visual Atlas 426

Gravity Visual Atlas 441

Adic and Polyphase Visual Atlas 457

Projectors, Packets, and Transfer: Visual Atlas

A visual atlas of resolution, currents, frames, and limits

1 An exact decomposition is a way of asking questions

A projector atlas divides a state into orthogonal packets and an operator into the transfers between those packets. The division can be fine or coarse, fixed or moving, and finite or countable. The useful question is always specific: which information survives the chosen representation, and what additional hypothesis makes the next operation legitimate?

These twelve plates build a common language before the specialized free Dirac construction. They begin with finite coordinates, where reconstruction is exact. They then introduce currents, changing frames, density matrices and tensor products. The final plates show what changes for positive measurement effects and for an unbounded operator on infinitely many coordinates. A finite commutative obstruction explains why ordinary pointwise differences cannot supply every desired differential identity.

Conventions. States are columns. The inner product is conjugate-linear in its first argument. An evolution generator A obeys $\dot{\psi} = A\psi$. For Hamiltonian evolution, $A = -iH$ and $\hbar = 1$. The first current index is the destination:

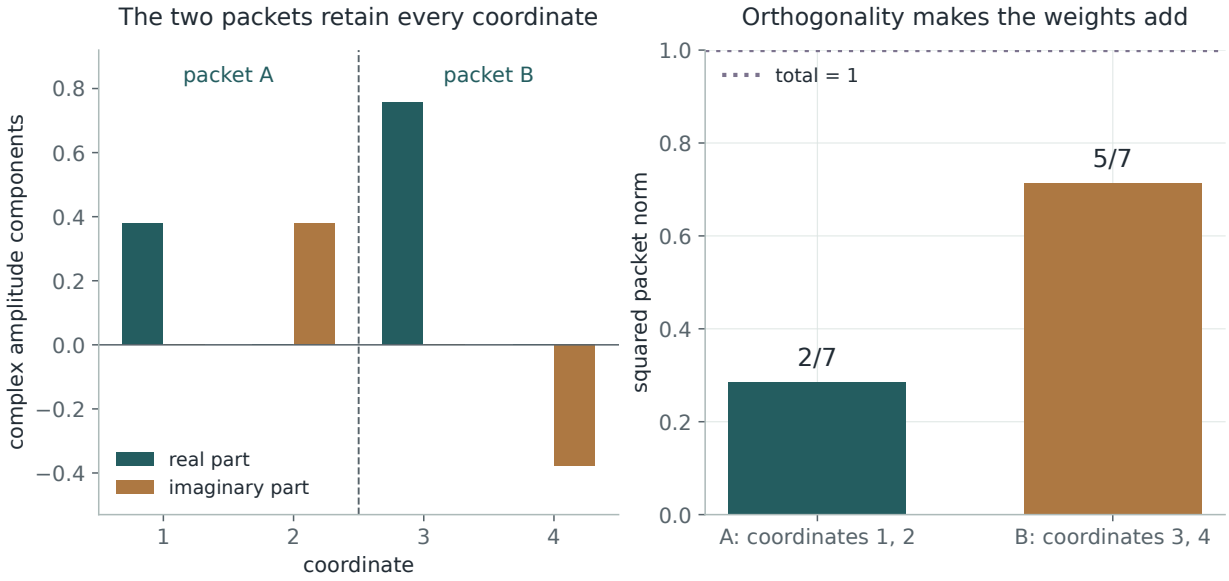
$$J_{a\leftarrow b} = 2\,\mathrm{Re}\,\langle P_a\psi, AP_b\psi\rangle = 2\,\mathrm{Im}\,\langle P_a\psi, HP_b\psi\rangle\,.$$

The diagonal infinite-dimensional examples use $He_n = ne_n$ for $n \geq 1$. They are transparent sequence models, rather than an identification of the full physical Hamiltonian with this particular spectrum.

What the pictures establish. Some plates display exact matrices. Others draw sampled values of explicit formulas. Their symbolic and numerical checks authenticate the examples and their implementation. The universal statements come from the arguments in the foundations monograph and the inherited source proofs. A curve that appears to level off does not prove that an infinite sum converges.

How the atlas connects to the formal work. The early Lean modules supply precise versions of the finite algebra, transfer identities, refinements, moving frames, density-matrix statements and positive-effect constructions. The companion Lean guide identifies exact declarations and hypotheses. The later diagonal examples explain analytic distinctions needed by the physics program. The atlas illustrates the mathematics and does not supply an additional compiler certificate.

2 Packets retain amplitudes, while weights retain less



An exact four-coordinate state. The left panel preserves its real and imaginary amplitudes. The right panel keeps only the two coarse packet weights. The weights add to one, but do not determine the original state.

Take

$$\psi = \frac{1}{\sqrt{7}}(1, i, 2, -i)^T, \qquad P_A = \mathrm{diag}(1, 1, 0, 0), \qquad P_B = I - P_A.$$

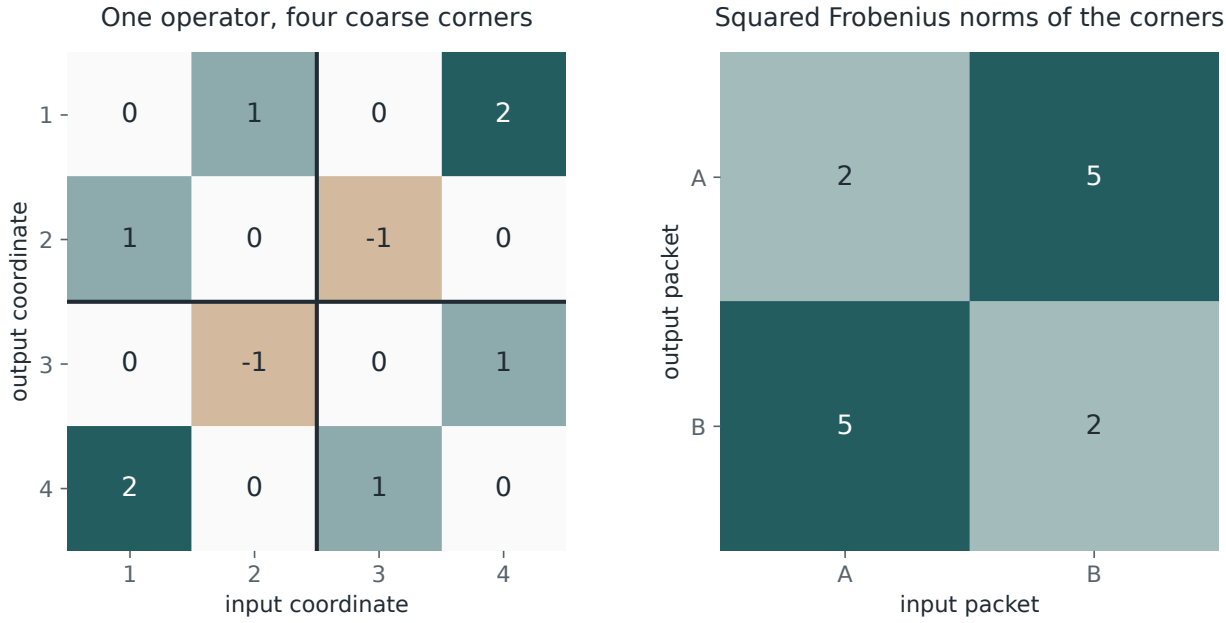
These are orthogonal projectors, so $P_A P_B = 0$ and $P_A + P_B = I$. The state is reconstructed exactly as $\psi = P_A\psi + P_B\psi$. Its squared norm splits without a cross term:

$$\|\psi\|^2 = \|P_A\psi\|^2 + \|P_B\psi\|^2 = \frac{2}{7} + \frac{5}{7} = 1.$$

The packets remain vectors in the original space. Replacing each packet with its squared norm loses information. For example, changing the sign of the fourth imaginary coordinate preserves both weights, while changing the state. Relative phases can later change interference and transfer.

Refinement separates coordinates that a coarse packet groups together. Coarsening sums orthogonal projectors and their populations. Neither operation requires discarding state information when all packet vectors are retained. An approximation begins only when some data are actually omitted or replaced.

3 An operator has an input packet and an output packet



Left: an exact Hermitian matrix, with the coarse division after coordinate 2. Right: squared Frobenius norms of its four corners. These are descriptive matrix sizes, not transition probabilities or packet populations.

For the same two projectors, every operator has four corners:

$$H = P_A H P_A + P_A H P_B + P_B H P_A + P_B H P_B.$$

The rightmost projector accepts the input packet and the leftmost projector selects the output packet. Thus $P_A H P_B$ carries the part of H that sends packet B into packet A. This convention fixes the orientation of every later transfer formula.

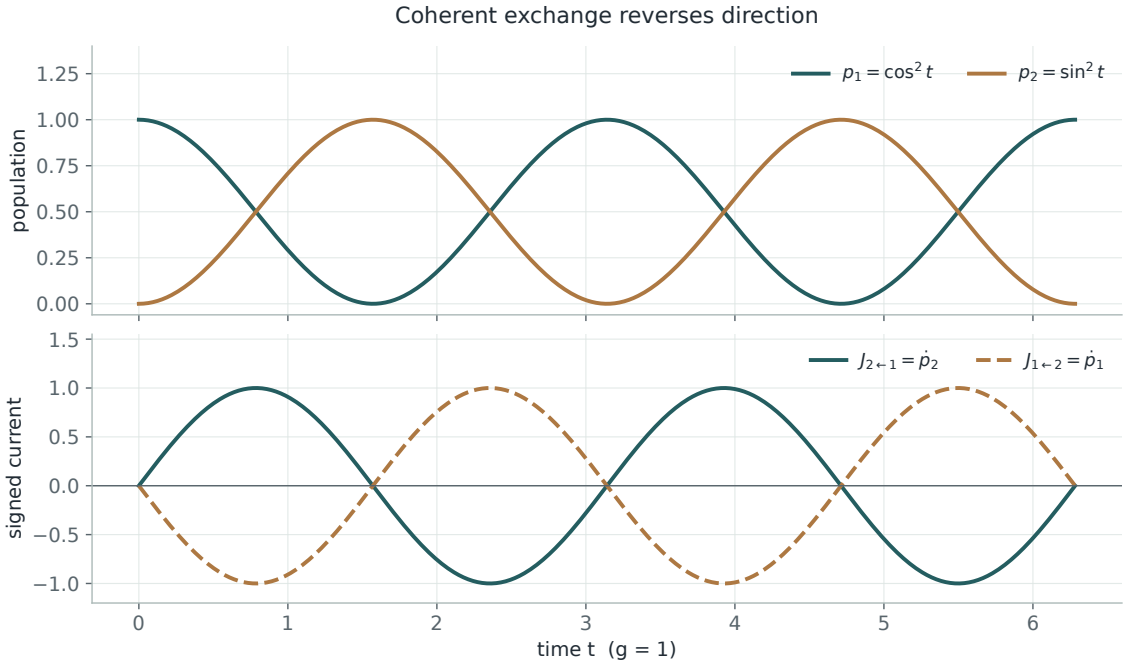
In the displayed example, the two within-packet blocks have squared Frobenius norm 2, and the two between-packet blocks have squared Frobenius norm 5. Their magnitudes alone do not determine a current. A current also depends on the evolving state and the phases in the relevant bilinear expression.

Composition introduces intermediate packets. Inserting the identity between two operators gives

$$P_a X Y P_b = \sum_c (P_a X P_c)(P_c Y P_b).$$

The sum records all two-step paths from the input packet b through c to the output packet a. This is exact matrix multiplication expressed through an atlas. Dropping an intermediate packet changes the operator unless a separate argument shows that its contribution vanishes or is acceptably small.

4 A current is signed, and its direction can reverse



A two-level Hamiltonian with coupling $g = 1$. The population curves exchange weight periodically. The lower panel uses destination-first currents, so a positive $J_{2\leftarrow 1}$ increases the second population.

Let $H = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$ and start from $\psi(0) = (1, 0)^T$. The exact solution is

$$\psi(t) = (\cos t, -i \sin t)^T, \qquad p_1(t) = \cos^2 t, \quad p_2(t) = \sin^2 t.$$

The current and its reverse are

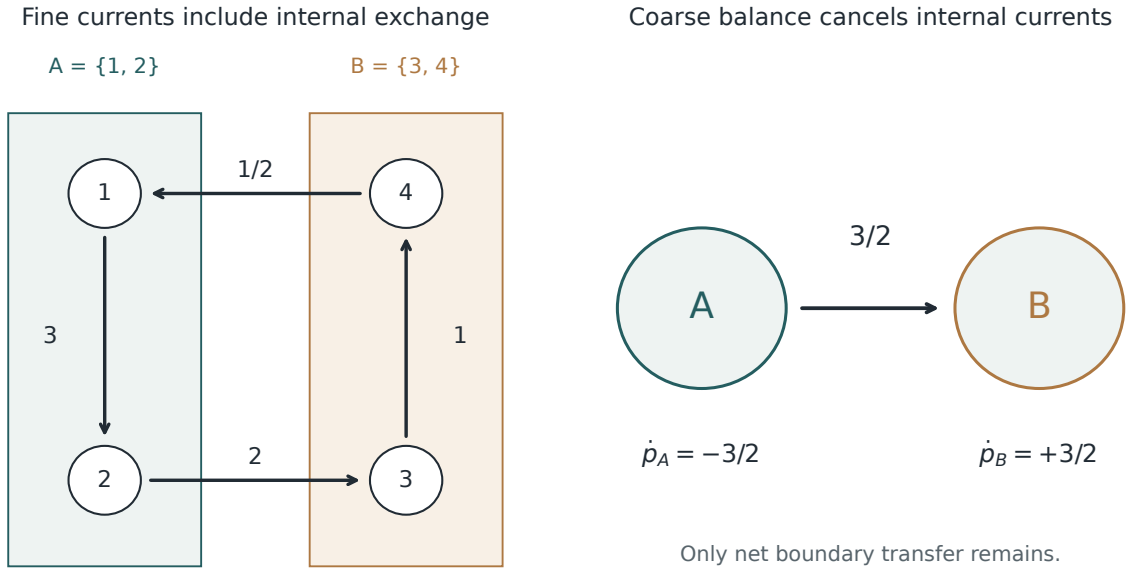
$$J_{2\leftarrow 1} = 2 \operatorname{Im}(\overline{\psi_2} \psi_1) = \sin(2t), \qquad J_{1\leftarrow 2} = -J_{2\leftarrow 1}.$$

Consequently $\dot{p}_2 = J_{2\leftarrow 1}$ and $\dot{p}_1 = -J_{2\leftarrow 1}$. Their sum is zero at every time. At a turning point of a population curve, the net current vanishes even though the coupling matrix is unchanged.

The sign distinguishes a current from a nonnegative transition rate. An antisymmetric current is a balance quantity. Its negative value means that the net transfer goes in the reverse direction. No interpretation as a classical stochastic rate follows from this identity alone.

The displayed solution is compared with an independent matrix exponential at 25 times. Its current is checked at 601 times. These checks catch sign and implementation errors. The exact formula supplies the all-time conclusion.

5 Coarsening cancels internal transfers



An abstract instantaneous antisymmetric current array on four packets. Positive arrows show the listed current orientation. Coarsening into $A = \{1, 2\}$ and $B = \{3, 4\}$ retains the net cross-boundary transfer $2 - 1/2 = 3/2$.

The fine currents in the drawing are $J_{2\leftarrow 1} = 3$, $J_{3\leftarrow 2} = 2$, $J_{4\leftarrow 3} = 1$ and $J_{1\leftarrow 4} = 1/2$. Their reverses have the opposite sign. The fine population derivatives are

$$\dot{p} = (-5/2,\ 1,\ 1,\ 1/2)^T.$$

Adding the first two entries gives $\dot{p}_A = -3/2$. Adding the last two gives $\dot{p}_B = 3/2$. The internal transfers of sizes 3 and 1 disappear from these coarse derivatives because each contributes once with each sign.

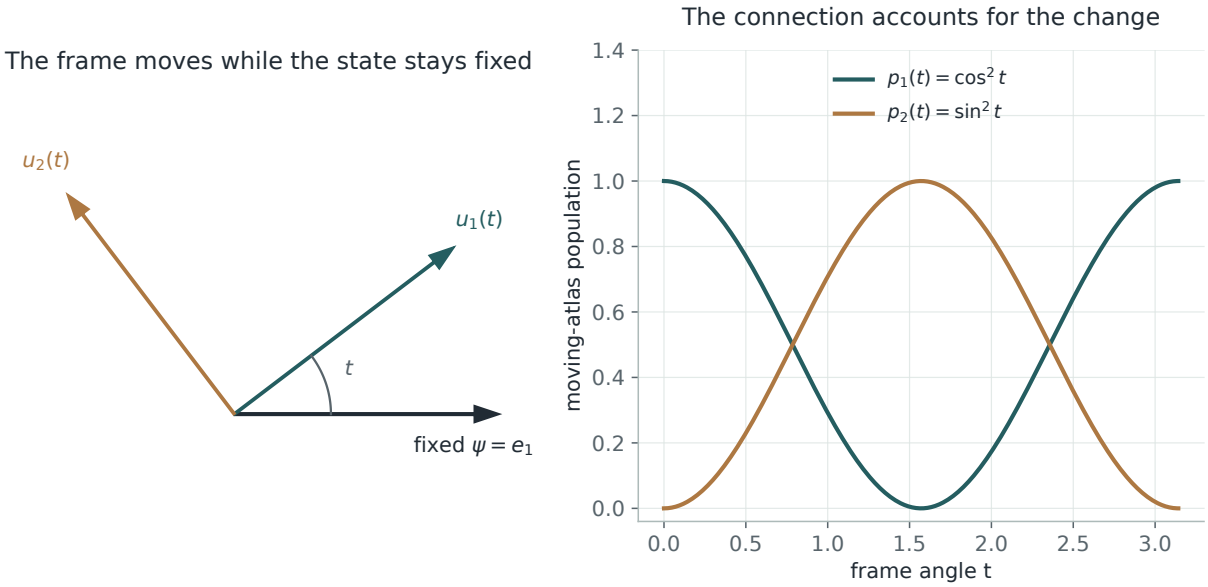
If R is the matrix that sums fine populations into coarse cells, the identities are

$$p^{\mathrm{c}} = Rp, \qquad J^{\mathrm{c}} = RJR^T, \qquad \dot{p}^{\mathrm{c}} = J^{\mathrm{c}}\mathbf{1}.$$

This is exact aggregation. The array is a balance example and is not presented as a constructed Hamiltonian trajectory.

Continuity does not determine every edge current. Adding a circulating current $1 \to 2 \to 3 \to 1$ of the same strength on all three edges changes no vertex divergence. In matrix notation, $K^T = -K$ and $K\mathbf{1} = 0$ imply $(J+K)\mathbf{1} = J\mathbf{1}$. A generator-induced current contains more information than population derivatives alone.

6 A moving atlas contributes its own connection



The laboratory state stays fixed at e_1, while the orthonormal frame rotates by the angle t. Its moving-atlas populations vary. The drawing on the left shows a single instant, and the curves on the right show the full angle interval.

Write

$$U(t) = \begin{pmatrix} \cos t & -\sin t \\ \sin t & \cos t \end{pmatrix}, \qquad \psi(t) = e_1, \qquad \phi(t) = U(t)^* \psi(t).$$

The moving coordinates are $\phi = (\cos t, -\sin t)^T$. Their squared magnitudes are the two plotted populations, even though the laboratory generator is $A = 0$.

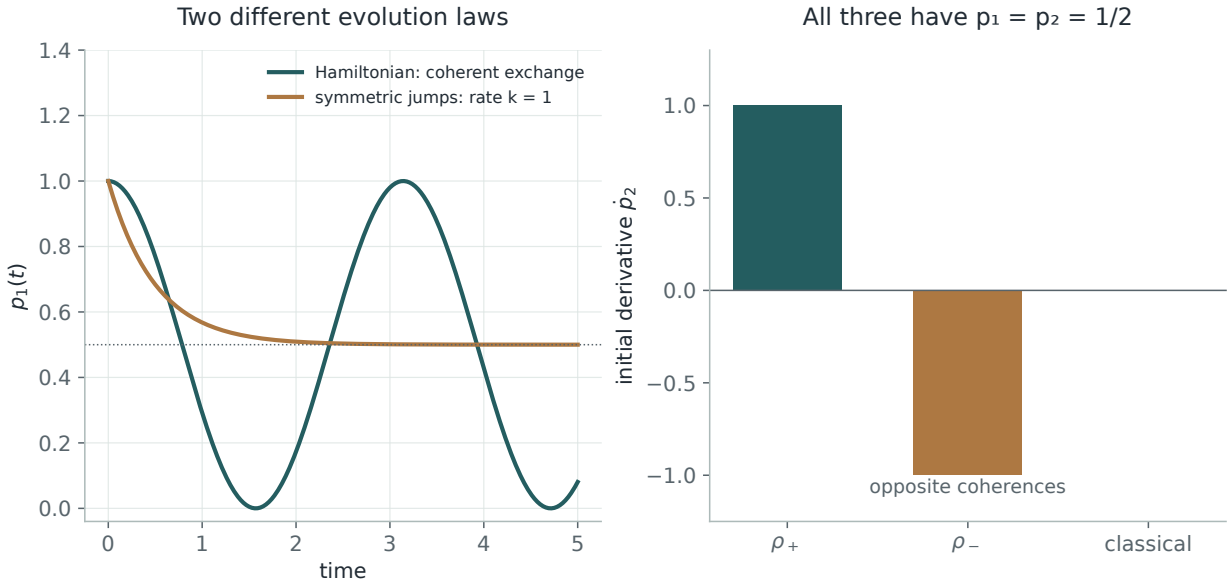
Differentiating the coordinate change gives

$$\Gamma = U^* \dot{U} = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \qquad \dot{\phi} = -\Gamma \phi.$$

For a general laboratory generator, this becomes $\dot{\phi} = (U^* A U - \Gamma)\phi$. The minus sign is fixed by differentiating $U^* U = I$. Both terms are needed to describe what a moving observer records.

This example separates a physical state change from a change in the measuring atlas. They can produce similar population curves while having different causes. Omitting the connection would predict constant moving coordinates for this stationary laboratory state, contradicting the explicit rotation.

7 Coherences and jump rates carry different information



Left: coherent Hamiltonian exchange and symmetric Lindblad jumps, both starting from population $(1, 0)$. Right: two pure density matrices and a classical equal-population state have the same populations but different initial derivatives under their stated evolution laws.

For $H = \left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$, the density equation is $\dot{\rho} = -i[H, \rho]$. The two pure states

$$\rho_+ = \frac{1}{2}\begin{pmatrix} 1 & i \\ -i & 1 \end{pmatrix}, \qquad \rho_- = \frac{1}{2}\begin{pmatrix} 1 & -i \\ i & 1 \end{pmatrix}$$

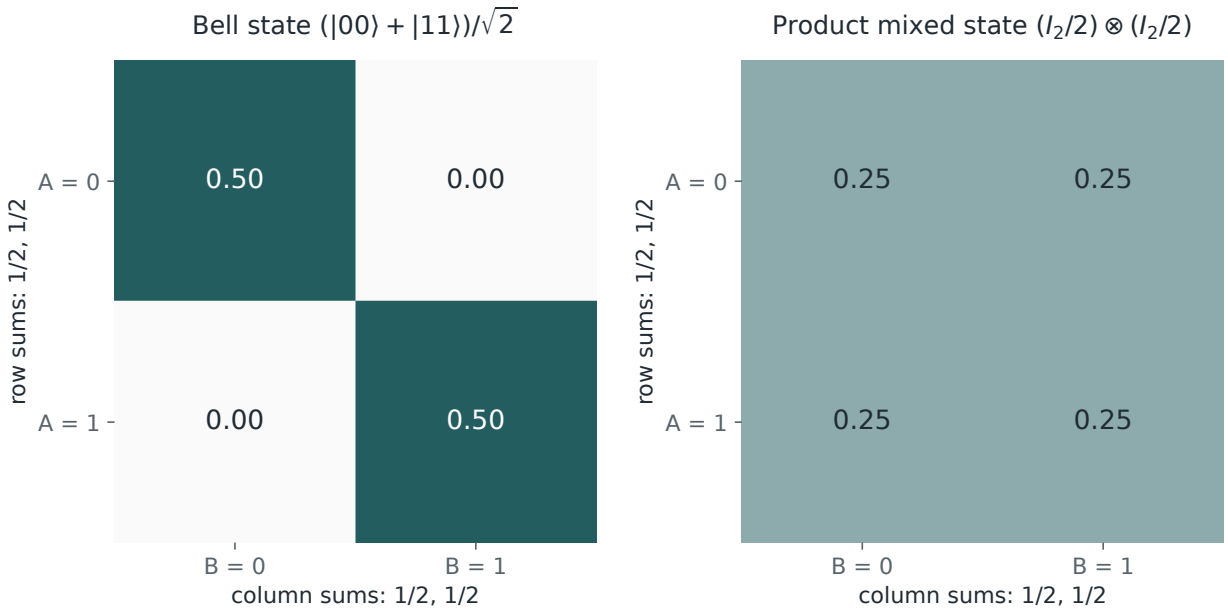
have the same populations $(1/2, 1/2)$. Yet they give $\dot{p}_2 = 1$ and $\dot{p}_2 = -1$, respectively. Their opposite off-diagonal coherences determine opposite initial transfers. A law written only in terms of these two populations cannot reproduce both derivatives.

For a different evolution, take the Lindblad jumps $L_{12} = |1\rangle\langle 2|$ and $L_{21} = |2\rangle\langle 1|$, with zero Hamiltonian. On diagonal density matrices they give

$$\dot{p}_1 = p_2 - p_1, \qquad p_1(t) = \tfrac{1}{2} + (p_1(0) - \tfrac{1}{2})e^{-2t}.$$

The left panel therefore contrasts periodic coherent exchange with relaxation generated by these specified jumps. This does not say that all Lindblad systems reduce to a classical two-state chain. Closure holds here because of the chosen jump operators and invariant diagonal family.

8 Marginals forget correlations



The computational-basis joint populations of a Bell state and a product mixed state. Each row and each column sums to $1/2$ in both panels. The joint tables differ even though both subsystem density matrices are $I_2/2$.

On $\mathbb{C}^2 \otimes \mathbb{C}^2$, let $\Phi = (|00\rangle + |11\rangle)/\sqrt{2}$. Its joint probabilities are $p_{00} = p_{11} = 1/2$ and $p_{01} = p_{10} = 0$. The product mixed state $(I_2/2) \otimes (I_2/2)$ instead gives all four outcomes probability $1/4$.

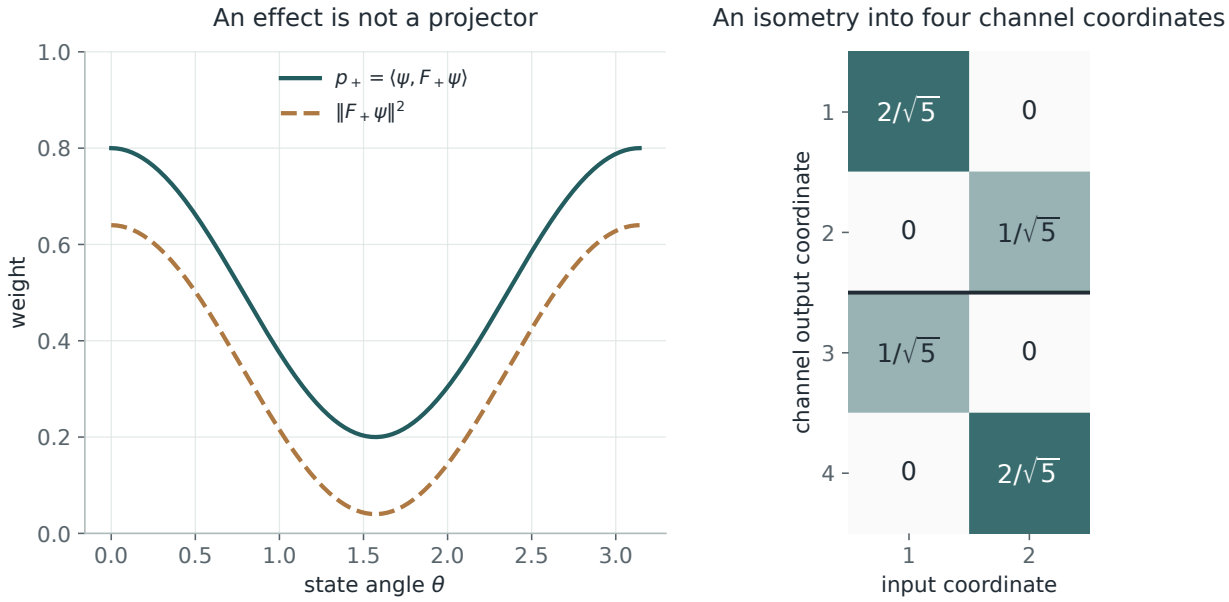
For tensor projectors $P_a \otimes Q_b$, the marginal identities follow by summing one resolution of the identity:

$$\sum_b p_{ab} = \mathrm{Tr}(\rho(P_a \otimes I)), \qquad \sum_a p_{ab} = \mathrm{Tr}(\rho(I \otimes Q_b)).$$

These operations recover subsystem populations from the joint state. The reverse direction is underdetermined. Even knowing both complete reduced density matrices does not determine the full joint density matrix.

The Bell state in this example is entangled, but this single joint probability table does not certify entanglement. The separable mixed state $\frac{1}{2}|00\rangle\langle 00| + \frac{1}{2}|11\rangle\langle 11|$ has the same table. One must retain or test additional information to distinguish its missing coherences from those of the Bell state.

9 Positive effects need channel amplitudes



An unsharp two-outcome measurement with $F_+ = \mathrm{diag}(4/5, 1/5)$ and $F_- = I - F_+$. Left: the probability $\langle \psi, F_+ \psi \rangle$ differs from $\|F_+ \psi\|^2$. Right: an explicit isometry realizes the probabilities as orthogonal channel norms in a larger space.

For $\psi_\theta = (\cos\theta, \sin\theta)^T$, the positive-effect probability is

$$p_+ = \frac{4}{5}\cos^2\theta + \frac{1}{5}\sin^2\theta.$$

Applying F_+ as though it were a projector instead gives the squared norm $\frac{16}{25}\cos^2\theta + \frac{1}{25}\sin^2\theta$. At $\theta = 0$, these are $4/5$ and $16/25$. The discrepancy is structural: $F_+^2 \neq F_+$.

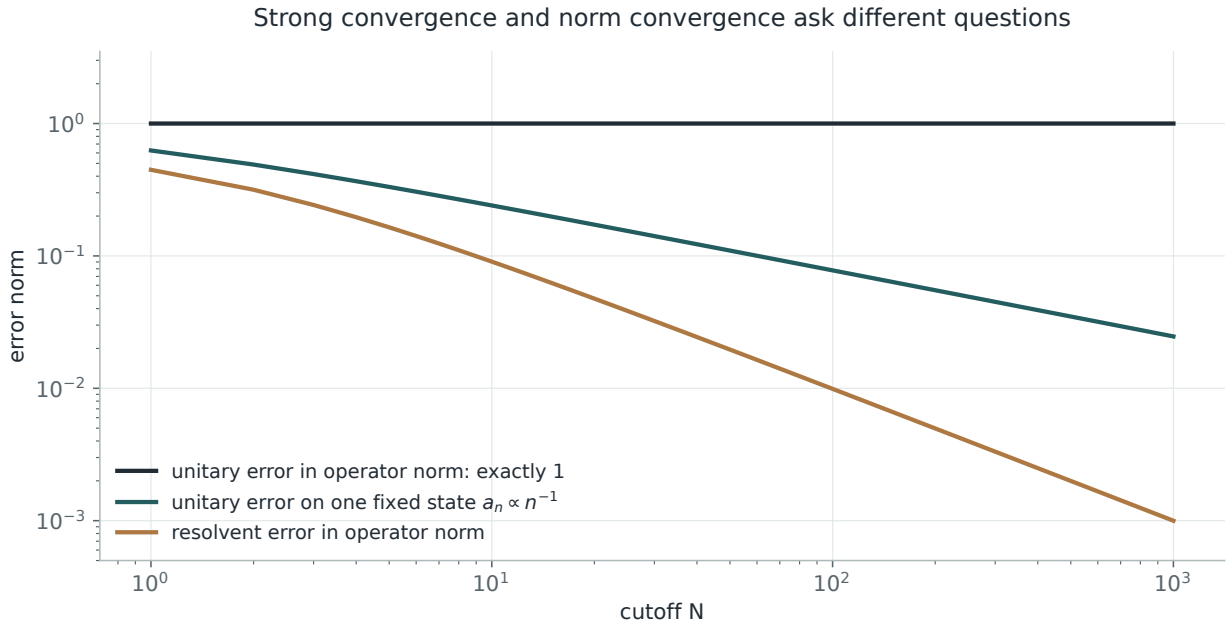
Use $V_+ = F_+^{1/2}$ and $V_- = F_-^{1/2}$ instead. Stacking the two channels gives an isometry $V : \mathbb{C}^2 \to \mathbb{C}^4$ with

$$V^*V = I, \qquad F_+ = V^*\Pi_+ V, \qquad p_+ = \left\|\Pi_+ V\psi\right\|^2,$$

where Π_+ selects the first two output coordinates. The channels are orthogonal in the enlarged space even though the original effects overlap: $F_+F_- \neq 0$.

This is an exact realization of a positive partition. It preserves probabilities and reconstruction through the analysis and synthesis maps. It does not make each original effect an idempotent projector.

10 A cutoff can converge strongly without converging in norm



For $He_n = ne_n$, retain the first N coordinates and set the omitted output to zero. The unitary error on one fixed state decays, while its operator norm remains exactly one. The resolvent has a decaying tail and does converge in operator norm.

Let P_N select coordinates $1, \dots, N$, and let $U(t)e_n = e^{-itn}e_n$. The zero-padded approximation is $P_N U(t) P_N$. Every omitted basis vector gives

$$\|(U(t) - P_N U(t) P_N)e_n\| = 1 \qquad (n > N).$$

The error has norm at most one because it is unitary evolution restricted to the tail, so its operator norm is exactly one for every finite N. The vector that witnesses this can move farther into the tail as N increases.

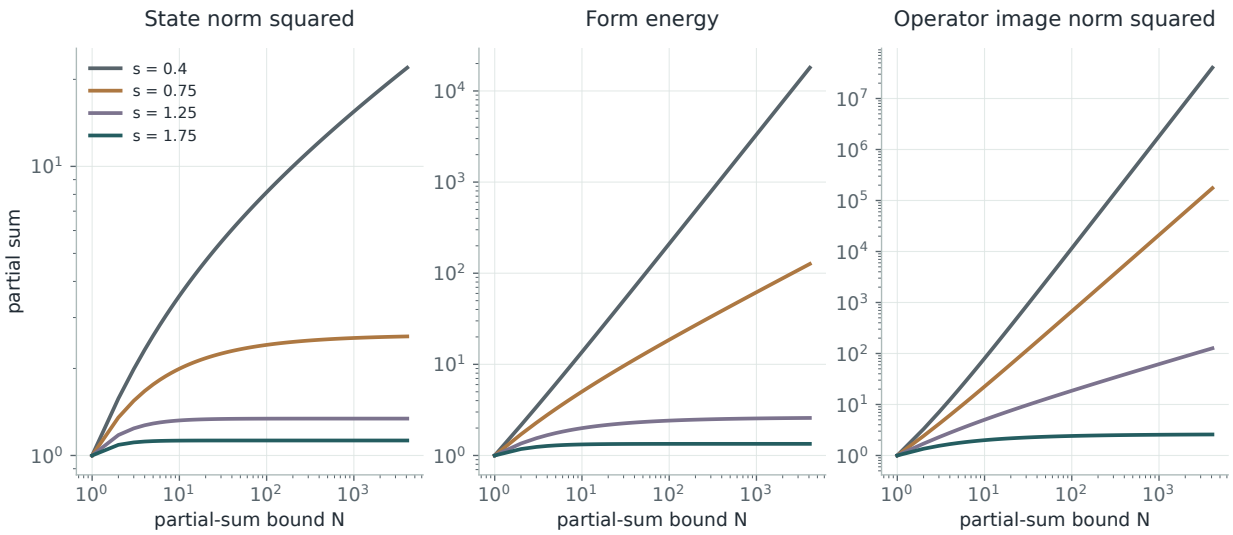
For the fixed normalized vector $a_n = n^{-1}/\sqrt{\zeta(2)}$, the error is instead

$$\|(U(t) - P_N U(t) P_N)a\| = \sqrt{\frac{\zeta(2, N+1)}{\zeta(2)}} \longrightarrow 0.$$

Here $\zeta(s, q)$ denotes the Hurwitz zeta tail. The error is independent of time in this commuting example. Strong convergence fixes a state first, whereas operator norm takes a supremum over all unit vectors.

The resolvent $(H - i)^{-1}$ has multipliers $(n - i)^{-1}$. Its cutoff error has operator norm $((N+1)^2 + 1)^{-1/2}$. Decay of the multiplier, rather than the mere existence of a cutoff, changes the norm-convergence conclusion.

11 Three spaces impose three different summability tests



Partial sums through $N = 4096$ for $a_n = n^{-s}$ and $He_n = ne_n$. The three panels weight the same coefficients by $1, n$ and n^2. Four exponents illustrate the distinct membership regimes. These finite plots do not decide convergence.

The exact tests are

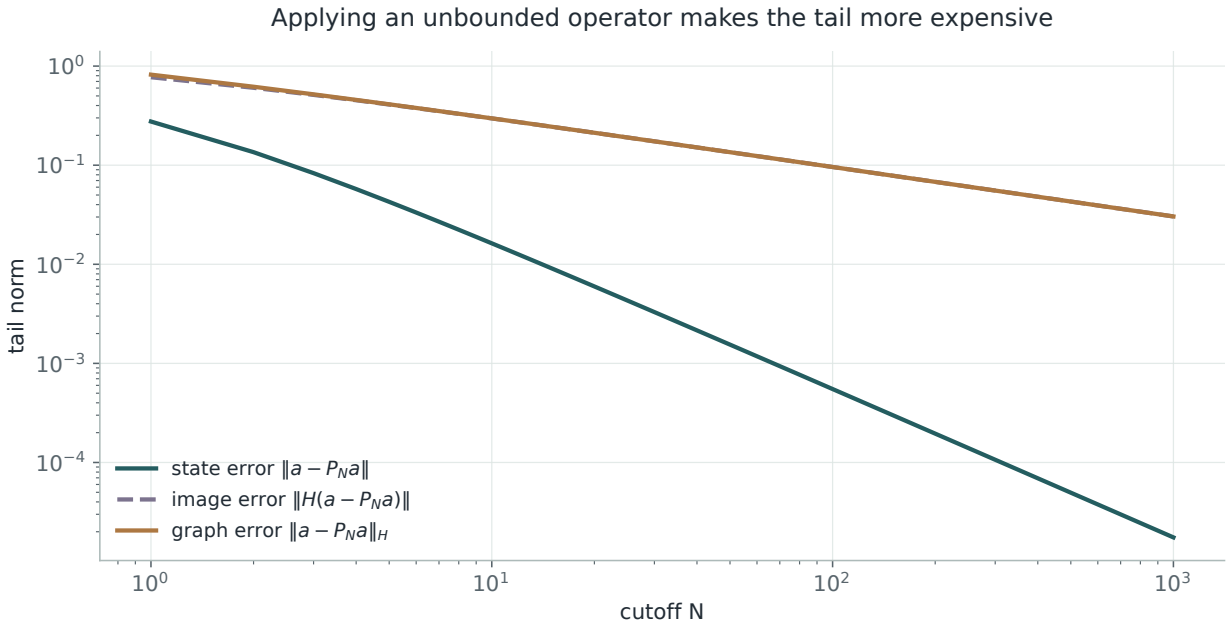
$$\begin{aligned} a \in \ell^2 &\iff \textstyle\sum_n n^{-2s} < \infty \iff s > \tfrac{1}{2}, \\ a \in D(H^{1/2}) &\iff \textstyle\sum_n n^{1-2s} < \infty \iff s > 1, \\ a \in D(H) &\iff \textstyle\sum_n n^{2-2s} < \infty \iff s > \tfrac{3}{2}. \end{aligned}$$

They follow from the integral test. Equality at any threshold produces a harmonic series and still fails the relevant test. The thresholds are strict.

Thus $s = 0.4$ does not define a Hilbert-space state. The choice $s = 0.75$ defines a state outside the form domain. The choice $s = 1.25$ belongs to the form domain but not to the operator domain. The choice $s = 1.75$ belongs to all three. These distinctions survive multiplication by a nonzero normalization constant.

For a positive unbounded operator, finite quadratic-form energy and an existing Hilbert-space image Ha are different requirements. A probability-normalized state need not satisfy either. Specifying the ambient Hilbert space alone therefore does not specify the domain of a proposed generator.

12 Graph convergence follows the state and its image



The normalized domain vector $a_n = n^{-2}/\sqrt{\zeta(4)}$ for $He_n = ne_n$. Its state tail falls much faster than its image tail. The graph curve combines both errors and lies just above the image curve.

The graph norm remembers what the operator does:

$$\|a\|_H^2 = \|a\|^2 + \|Ha\|^2 .$$

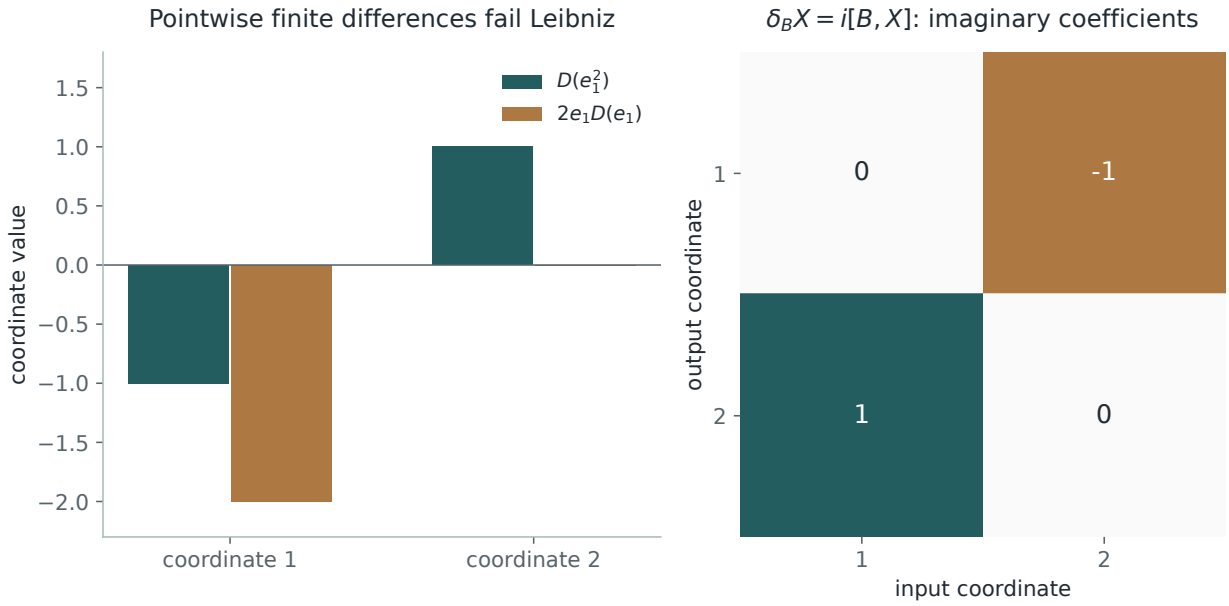
For the displayed vector, each relevant infinite tail is explicit:

$$\begin{aligned} \|a - P_N a\|^2 &= \frac{\zeta(4, N+1)}{\zeta(4)}, \\ \|H(a - P_N a)\|^2 &= \frac{\zeta(2, N+1)}{\zeta(4)}, \\ \|a - P_N a\|_H^2 &= \frac{\zeta(4, N+1) + \zeta(2, N+1)}{\zeta(4)}. \end{aligned}$$

The ordinary state error is of order $N^{-3/2}$. The image and graph errors are of order $N^{-1/2}$. A visually small omitted state tail can therefore carry a substantially larger operator error.

Here the cutoff commutes with H and preserves its domain, so the weighted tail proves graph convergence. For a general countable atlas, those properties need hypotheses. Strong reconstruction of the identity on the ambient Hilbert space alone does not establish graph compatibility with an unbounded operator.

13 Exact finite differentiation requires the right algebra



Left: a nonzero two-point difference violates the ordinary pointwise Leibniz identity on an idempotent. Right: the nonzero matrix derivation $i[\sigma_x, \mathrm{diag}(1,0)]$. The matrix panel displays coefficients of i, so its off-diagonal entries are $-i$ and i.

In the finite pointwise algebra $\mathbb{C}^N$, each coordinate indicator e_a satisfies $e_a^2 = e_a$. An exact linear derivation would have to obey

$$De_a = D(e_a^2) = 2e_a De_a.$$

Outside coordinate a, this forces every component to zero. At coordinate a, it says $x = 2x$, so that component is also zero. The indicators span $\mathbb{C}^N$, hence every such derivation is zero.

The left panel uses $D(f_1, f_2) = (f_2 - f_1, f_1 - f_2)$. It is a useful finite difference, but $D(e_1^2) = (-1, 1)$ while $2e_1De_1 = (-2, 0)$. The obstruction concerns the exact pointwise product rule, not the usefulness of discrete approximations.

In $M_N(\mathbb{C})$, a Hermitian B instead gives the derivation $\delta_B X = i[B, X]$. Expanding the commutator proves

$$\delta_B(XY) = (\delta_B X)Y + X\delta_B Y, \qquad \mathrm{Tr}((\delta_B X)Y) = -\mathrm{Tr}(X\delta_B Y).$$

The algebra has changed, and its additional corners carry nontrivial derivatives. The finite example uses $B = \sigma_x$ and checks the two identities on all 16 pairs of 2×2 matrix units.

14 What the twelve plates preserve, and what they forget

The atlas is organized around a sequence of exact distinctions. Packet vectors reconstruct a state, while packet weights omit phase information. Corners reconstruct an operator, while a corner norm omits the bilinear dependence of a current on a state. Coarse populations preserve net boundary balance, while they forget internal currents and divergence-free cycles.

A moving atlas introduces another source of population change. Its connection records the coordinate motion. A density matrix retains the coherences that a population vector alone cannot determine. Tensor marginals retain local information while discarding some joint correlations. Positive effects retain normalized measurement probabilities, but their original-space images are not orthogonal packets.

The infinite-dimensional passage makes the topology and the domain part of the statement. A cutoff that approximates every fixed state may still have operator-norm error one. A vector in the ambient Hilbert space may lie outside the form domain or the operator domain. A graph approximation must control the image under the operator as well as the vector itself.

These are useful restrictions. They identify the exact data and hypotheses that the next result must supply. The free Dirac monograph uses the same distinctions when it passes from finite matrix actions to full coefficient evolution, Fourier transport, an unbounded Hamiltonian and its weighted domain. Keeping these passages separate makes the operator construction and its hypotheses precise.

A practical reading order. For the finite calculus, read plates 1 through 5 together and then return to the transfer proofs. For probability and composite systems, compare plates 6 through 8. For the unbounded-operator frontier, read plates 9 through 11 before the domain arguments in the Dirac paper. Plate 12 explains why a desired differential identity cannot be assumed merely because a finite model has coordinate differences.

Reproduction and audit boundary. Run `python tools/make_foundation_figures.py` from the source root. It produces twelve PDF figures, twelve PNG previews and `evidence/foundations/figure_checks.json`. The manifest records exact inputs, software dependencies, finite ranges, tolerances, output hashes and the limitations of the calculations. No random samples are used. The floating computations are not interval certified, and no Lean compilation is performed by this command.

The figures are explanations of specified mathematical objects. Their checks authenticate implementation and finite examples. They are not a new empirical validation of the physics, and they do not turn an inherited conditional statement into an unconditional theorem.

Fourier Packets and the Continuum Dirac Operator

A visual atlas of modes, evolution, approximation and operator domains

1 How to read this atlas

The development moves from finite algebra to an infinite-dimensional operator. Each passage changes the question that must be answered. A matrix may preserve the norm of one spinor without defining an operator on every square-summable field. A finite Fourier calculation may be exact without controlling the omitted modes. A symmetric formula still needs a specified domain and an adjoint-domain argument before self-adjointness is established.

The following twelve plates isolate these transitions. They follow a single set of conventions: the unit torus $\mathbb{T}^3 = (\mathbb{R}/\mathbb{Z})^3$, normalized Haar measure, Fourier momentum $p(n) = 2\pi n$, and four-component column spinors with coordinates $0, 1, 2, 3$. Time evolution uses $U(t) = \exp(-itH)$ on each finite matrix. The centered symbol is

$$q_h(n)_j = \frac{\sin(2\pi h n_j)}{h} \quad (h \neq 0), \qquad q_0(n) = 2\pi n.$$

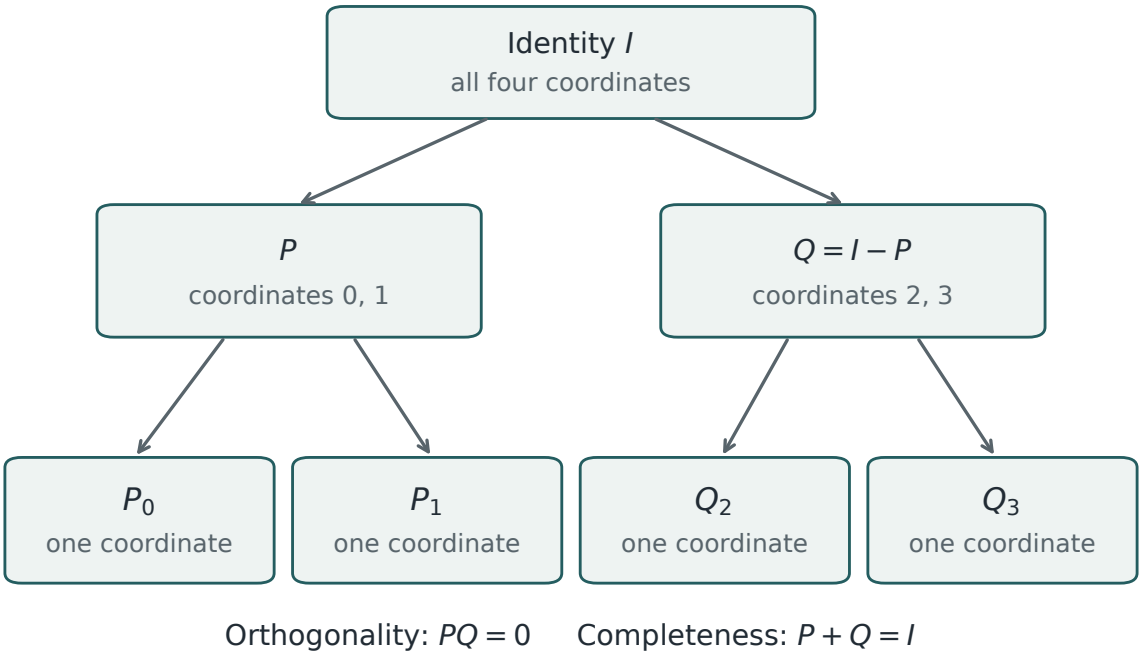
This is a Fourier multiplier on the same torus. The integer modes have not been identified modulo a sampled grid.

Three kinds of visual evidence. The projector tree and transport square explain definitions and relationships. Their geometry is schematic. The matrix display records exact entries. Dispersion, evolution, and tail plots are recomputed numerical illustrations of explicit formulas. A smooth curve does not certify a theorem, and a finite partial sum does not decide convergence of an infinite series. The accompanying text supplies the analytic distinction in each case.

The complete free-field argument. The operator construction supplies full coefficient and physical evolution, a self-adjoint Hamiltonian, finite Fourier graph approximation, and the exact first-order domain. The checked generator connects the group to this operator. The centered extension proves graph and strong-resolvent limits and the quadratic rate under a finite third-order Fourier moment. The abstract spectral-functional-calculus identification remains a separate written deduction.

Reproduction. The retained figure generator is `reference_record/atlas_tools/make_figures.py`. The script checks the Clifford relations in exact symbolic arithmetic and compares the displayed propagator with an independent matrix-exponential calculation on a stated finite set. The script specifies its software dependencies, tolerances and finite ranges, and writes a manifest named `evidence/figure_checks.json` in its original tree. That manifest is not present in this publication copy; these atlas computations were not rerun for the current formal extension. The current scope schematic is separately regenerated by `tools/render_formal_scope.py` from an accepted formal receipt. These checks validate the examples used by the atlas. They do not substitute for the Lean source or its verification receipts.

2 Refinement preserves an exact decomposition



A four-dimensional coordinate example. Each child is an orthogonal projector and the children sum to their parent. The drawing is a schematic of refinement, not an identification with spatial dyadic boxes.

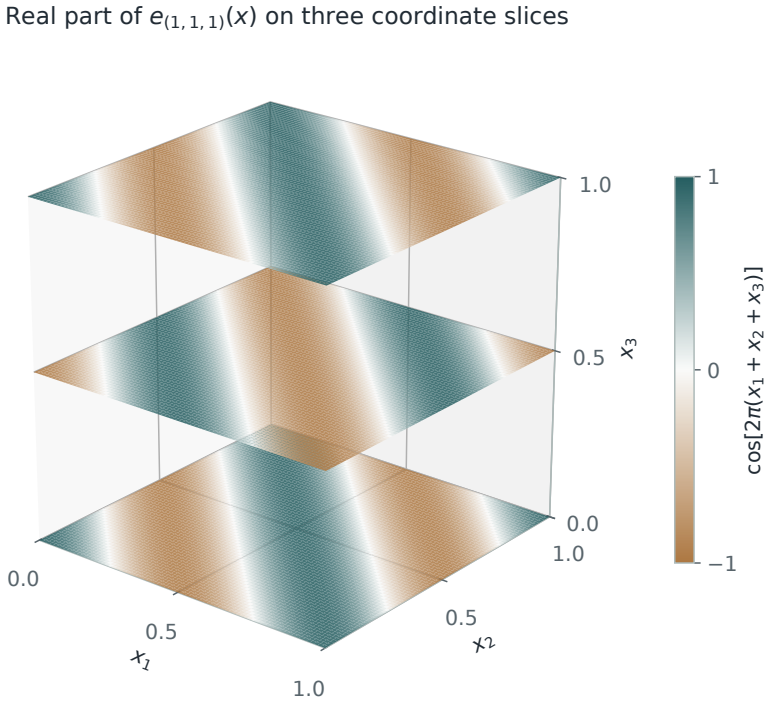
Take $P = \mathrm{diag}(1, 1, 0, 0)$ and $Q = I - P$. The first split distinguishes two orthogonal sectors. Refining P into its first two coordinate projectors preserves its range exactly. Refining Q does the same for the remaining coordinates:

$$P = P_0 + P_1, \qquad Q = Q_2 + Q_3, \qquad I = P_0 + P_1 + Q_2 + Q_3.$$

An operator T can be resolved into transfer blocks $P_a T P_b$. A block records which input sector is accepted and which output sector is retained. Summing all blocks recovers T when the projectors resolve the identity. Off-diagonal blocks describe transfer between distinct sectors, while diagonal blocks describe action within one sector.

The value of the language is its bookkeeping. It distinguishes an exact decomposition from an approximation that discards some blocks. This elementary example also explains why a change of representation must preserve both the projectors and the inner product when it is used to transport norm claims.

3 A Fourier mode lives on the three-torus



Three genuine three-dimensional slices of the real part of $e_{(1,1,1)}(x)$. Equal values on opposite faces reflect periodicity. Color shows a real component, not probability density.

The scalar mode $e_n(x) = \exp(2\pi i n \cdot x)$ has absolute value one everywhere. For normalized Haar measure, its squared L^2 norm is also one. Distinct integer modes are orthogonal:

$$\int_{\mathbb{T}^3} \overline{e_n(x)} e_k(x)\, dx = \delta_{nk}.$$

The plate shows the mode $(1, 1, 1)$ at $x_3 = 0, 1/2, 1$. Its real part changes sign, while the magnitude of the full complex wave stays fixed. The upper and lower faces have the same pattern because adding one to a coordinate changes the phase by an integer multiple of 2π.

A spinor amplitude $a_n \in \mathbb{C}^4$ multiplies this scalar mode. For a finite sum $f(x) = \sum_n e_n(x) a_n$, orthogonality removes the cross terms only after integration. It gives $\|f\|_{L^2}^2 = \sum_n \|a_n\|^2$. Local interference can still produce a nonconstant density. This distinction is essential when reading the evolution plate.

4 The coordinate convention determines the signs



Real and imaginary parts of the four supplied 4×4 generators. Blank cells are zero. The lower panel records coefficients of i, so a displayed -1 denotes the entry $-i$.

Rows and columns follow the spinor order $0, 1, 2, 3$. In particular,

$$\alpha_y v = (-iv_3,\ iv_2,\ -iv_1,\ iv_0)^T.$$

The sign in this formula fixes the worked packet's fourth component. Replacing it silently with a different convention would change the explicit evolution even if another Clifford representation remained mathematically valid.

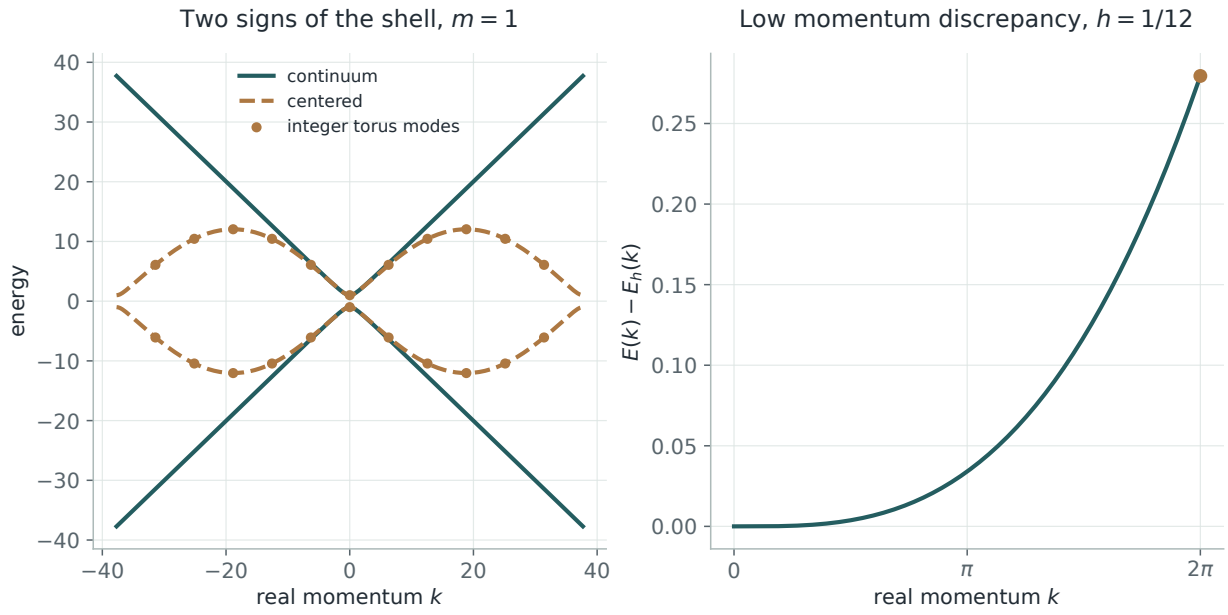
The matrices are Hermitian involutions and anticommute in distinct directions. Therefore the modal Hamiltonian

$$H(m,p) = m\beta + p_x\alpha_x + p_y\alpha_y + p_z\alpha_z$$

satisfies $H(m,p)^2 = (m^2 + |p|^2)I$. The cross terms vanish in pairs. This single identity explains the energy shell and the trigonometric formula for the matrix exponential. Hermiticity supplies exact norm preservation.

The figure generator checks all four squares and all six distinct anticommutators in exact Gaussian-integer arithmetic. That small computation is a reproducibility check of the displayed matrices, separate from the formal proofs in the project.

5 Dispersion changes when momentum becomes a sine



Continuum and centered axis energies at $m = 1$, $h = 1/12$. Curves interpolate real momentum. Dots mark actual integer torus modes on the positive and negative centered branches. The right panel enlarges the low momentum energy difference.

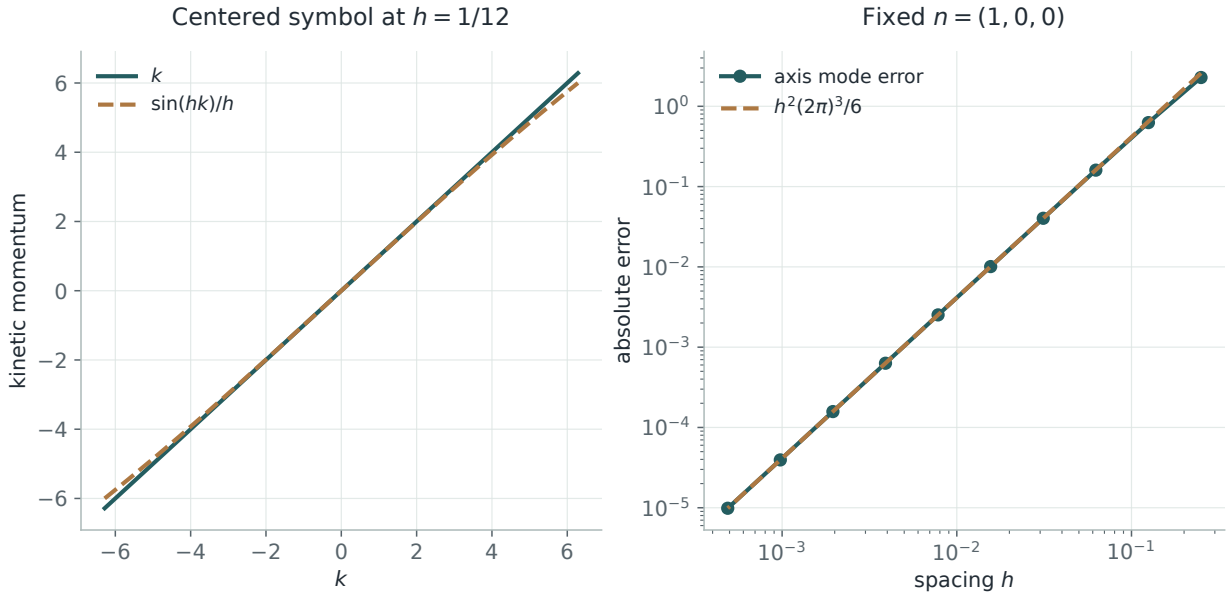
The square of the Hamiltonian is scalar, so its two shell magnitudes are

$$E(p) = \sqrt{m^2 + |p|^2}, \qquad E_h(p) = \sqrt{m^2 + \sum_j [\sin(hp_j)/h]^2}.$$

Each nonzero shell has positive and negative spectral sectors. The centered symbol is close to p when $|hp_j|$ is small. Over a larger real momentum interval it bends and returns, whereas the continuum shell continues to grow. These global curves should not be read as uniform approximation over all modes.

For the actual torus axis mode $n = (1, 0, 0)$, $p = (2\pi, 0, 0)$ and $q_{1/12} = (6, 0, 0)$. The corresponding positive energies are $\sqrt{1 + 4\pi^2}$ and $\sqrt{37}$. The torus mode labels remain integers throughout. A curve connecting those labels is a visual interpolation, not an additional family of torus eigenfunctions.

6 Quadratic consistency has a fixed-mode scope



The centered sine momentum approaches the continuum momentum. At the fixed axis mode $n = (1, 0, 0)$, the numerical error follows the analytic quadratic bound as h decreases. The sample spacings are $2^{-2}, \ldots, 2^{-11}$.

The elementary sine remainder gives

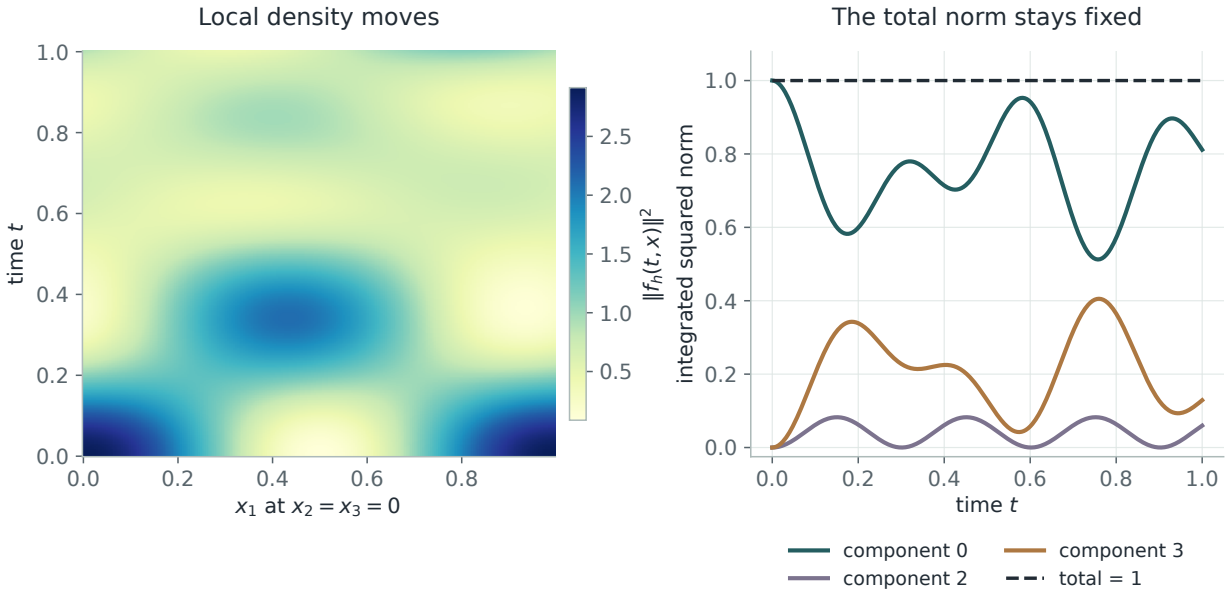
$$\left|\frac{\sin(hp_j)}{h} - p_j\right| \leq \frac{|h|^2|p_j|^3}{6}.$$

For a fixed finite set of modes, the momentum factors have a finite maximum. This yields a quadratic consistency estimate with a constant depending on that set. Exact Hermitian evolution then converts a modal Hamiltonian error into an evolution error growing at most linearly with $|t|$.

The dependence on momentum matters. The factor $|p_j|^3$ grows without bound as the integer mode increases. One cannot remove a Fourier cutoff and retain the same quadratic constant for every L^2 field. Additional weighted regularity is needed to sum a rate estimate of this form over infinitely many coefficients.

The checked full-space convergence result instead controls a chosen finite core and its remaining tail separately. The next two plates show why exact stability makes that passage possible without assigning a universal quadratic rate to arbitrary data.

7 Local density moves while total norm is conserved



The exact three-mode centered packet with $m = 1$, $h = 1/12$. Left: pointwise density on the line $x_2 = x_3 = 0$. Right: integrated component weights and their conserved total. A line through the torus is a display slice, not the integration domain.

Use modes $(0,0,0)$, $(1,0,0)$ and $(1,1,1)$, with initial spinors respectively $e_0/\sqrt{2}$, $e_0/2$ and $e_0/2$. Their coefficient norm squares sum to one. Each coefficient evolves by

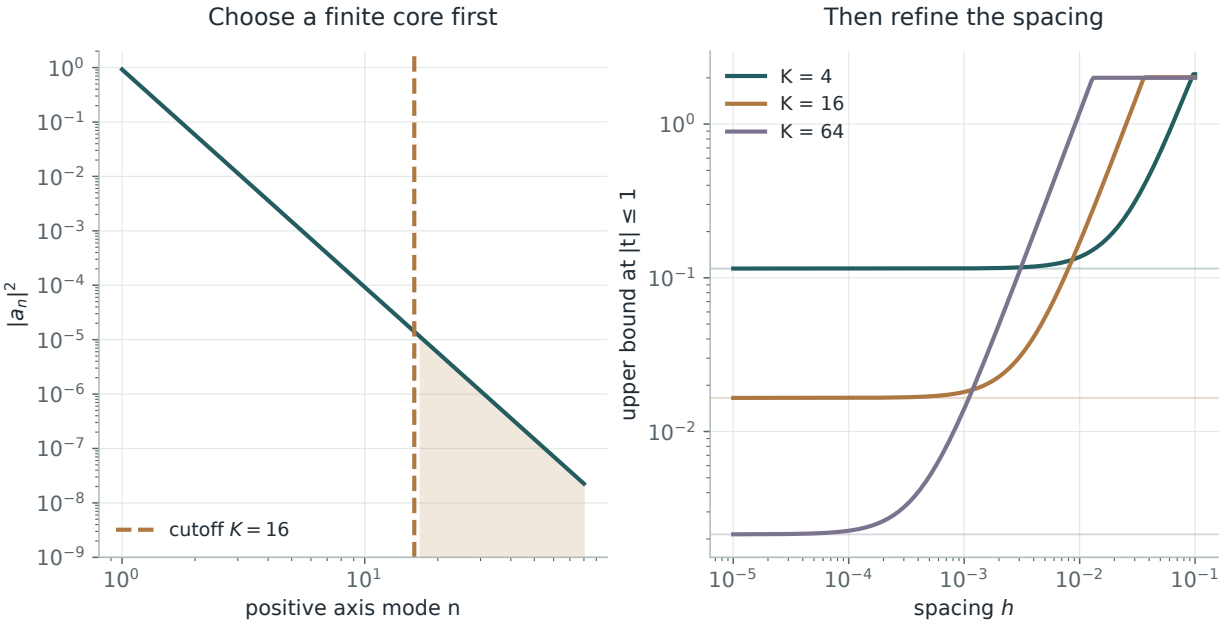
$$U_n(t) = \cos(E_n t) I - i \frac{\sin(E_n t)}{E_n} H_n,$$

with the continuous zero-energy branch interpreted as I.

The field is a coherent sum of these evolving modes. Cross terms redistribute the local density, and individual spinor components exchange weight. Fourier orthogonality and modal unitarity nevertheless give $\|f_h(t)\|_{L^2}^2 = 1$ for every real time. Component 1 remains zero in this particular example and is omitted from the right-hand legend.

The atlas checks the formula against an independent matrix exponential and verifies the total norm on 401 times in $[0,1]$. The all-time statement comes from the algebraic and Fourier arguments, not from the sampled residual.

8 Removing the cutoff requires two choices



An axis-supported example $a_n = n^{-2}/\sqrt{\zeta(4)}$ for $n \ge 1$. Left: coefficient mass and a selected cutoff. Right: explicit bounds at $|t| \le 1$ for three fixed cutoffs. The horizontal floors are the two-tail term.

Write $P_K a$ for a finite Fourier truncation. The two exact evolutions preserve the norm, so

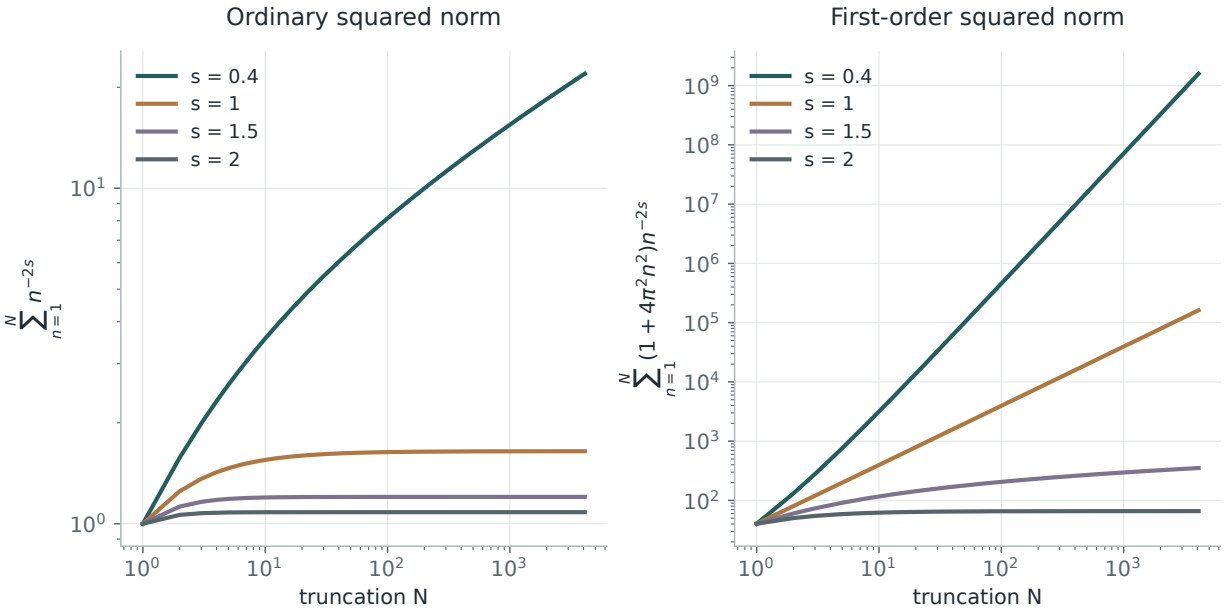
$$\|(U_h(t) - U(t))a\| \le \|(U_h(t) - U(t))P_K a\| + 2\,\|(I - P_K)a\|\,.$$

The correct order is to choose K so that the tail is small, then choose h so that the finite-core term is small uniformly over the time interval. The plot uses the smaller of the finite-core quadratic bound and the elementary stability bound $2\,\|P_K a\|$.

Holding K fixed leaves a positive tail floor even as $h \to 0$. Increasing K lowers that floor but increases the coefficient in the finite-core estimate. This tradeoff is not a failure of convergence. It explains why a density-and-stability proof establishes convergence without automatically providing one quadratic rate for all data.

The selected axis sequence is only an illustration. The checked density argument applies to arbitrary square-summable coefficient fields, with finite supports chosen for those fields.

9 Square summability is weaker than operator membership



Finite partial sums through $N = 4096$ for axis-supported coefficients $a_{(n,0,0)} = n^{-s}e_0$, $n \geq 1$. The zero mode and every other mode are zero. The exact classifications below use the power-series criterion, not the apparent flattening of a curve.

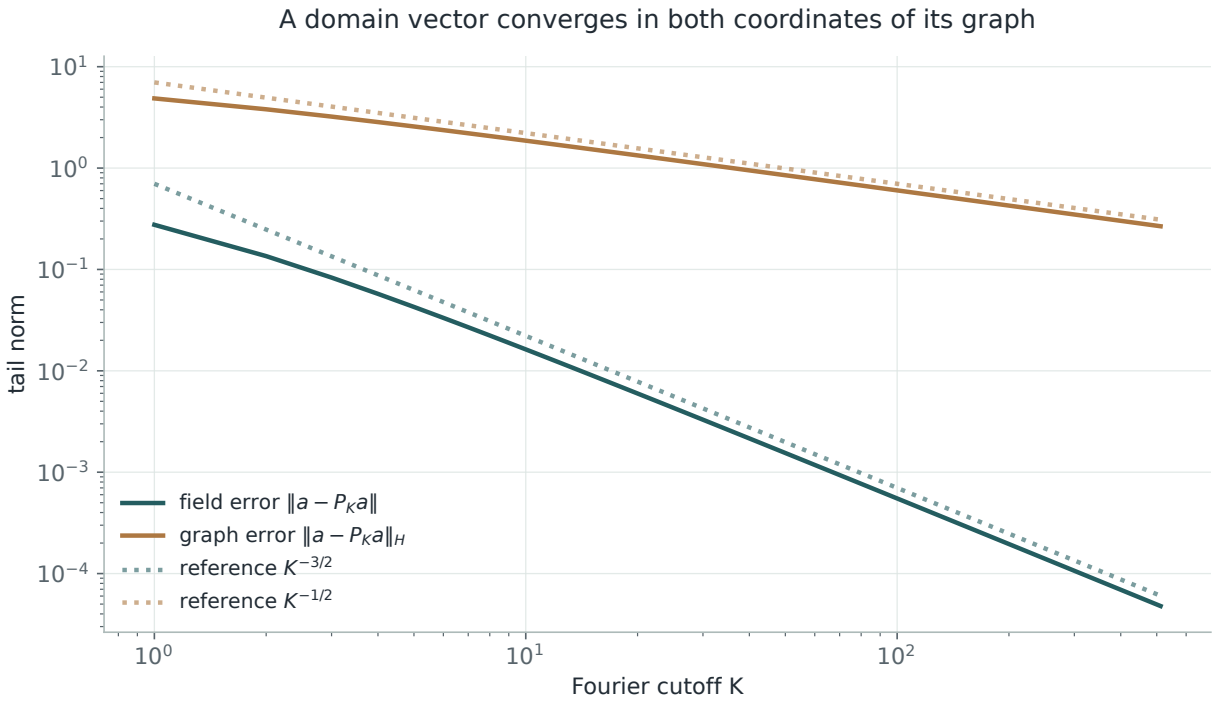
For this one-dimensional support inside $\mathbb{Z}^3$, ordinary square summability and first-order weighted summability have different thresholds:

$$\sum_{n\geq 1} n^{-2s} < \infty \iff s > \tfrac{1}{2}, \qquad \sum_{n\geq 1} (1 + 4\pi^2 n^2) n^{-2s} < \infty \iff s > \tfrac{3}{2}.$$

The choice $s = 0.4$ does not define an L^2 coefficient field at all. The choice $s = 1$ defines an L^2 field outside the Dirac domain. The borderline $s = 1.5$ still lies outside that domain because its derivative-energy series contains a harmonic tail. The choice $s = 2$ belongs to the domain.

These are thresholds for the stated axis-supported family. They must not be transferred unchanged to radially populated three-dimensional families, where the number of modes in each shell also grows. The operator domain is a condition on the complete weighted sum, not a visual smoothness judgment.

10 A graph approximation controls the field and its image



Truncation errors for the normalized positive-axis sequence $a_n = n^{-2}/\sqrt{\zeta(4)}$ at $m = 1$. Both curves use infinite tail formulas, evaluated with the Hurwitz zeta function. Dotted lines show asymptotic slopes rather than fitted bounds.

The graph norm is

$$\|a\|_{H_m}^2 = \|a\|^2 + \|H_m a\|^2 .$$

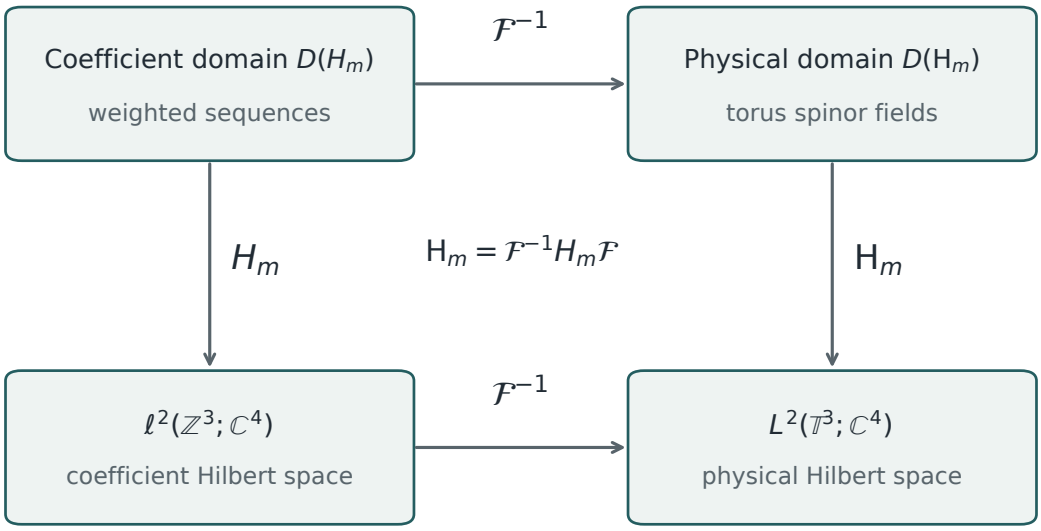
Convergence in ordinary ℓ^2 controls the field. Graph convergence also controls the result of applying the unbounded operator. For the displayed sequence,

$$\|a - P_K a\|^2 = \frac{\zeta(4, K+1)}{\zeta(4)}, \qquad \|a - P_K a\|_{H_1}^2 = \frac{2\zeta(4, K+1) + 4\pi^2 \zeta(2, K+1)}{\zeta(4)}.$$

The ordinary error decays like $K^{-3/2}$ while the graph error decays like $K^{-1/2}$. Both vanish, but the unbounded multiplier makes high modes more expensive in the second measurement.

The maximal-domain argument proves the structural statement behind this example: finite truncations preserve the domain, commute with the multiplier, and approximate a domain vector together with its image. This gives a finite Fourier graph core. No observed slope is needed for that existence statement.

11 Fourier transport carries the domain as well as the formula



The horizontal maps preserve inner products and norms.

The inverse Fourier equivalence carries coefficient fields to physical torus fields. The square commutes by the definition of the transported operator. The upper horizontal map acts on the exact domains, and the lower map acts on the ambient Hilbert spaces.

Let $\mathcal{F}$ send a physical field to its Fourier coefficients. The physical Hamiltonian is defined by

$$\mathsf{H}_m = \mathcal{F}^{-1} H_m \mathcal{F}, \qquad D(\mathsf{H}_m) = \mathcal{F}^{-1} D(H_m).$$

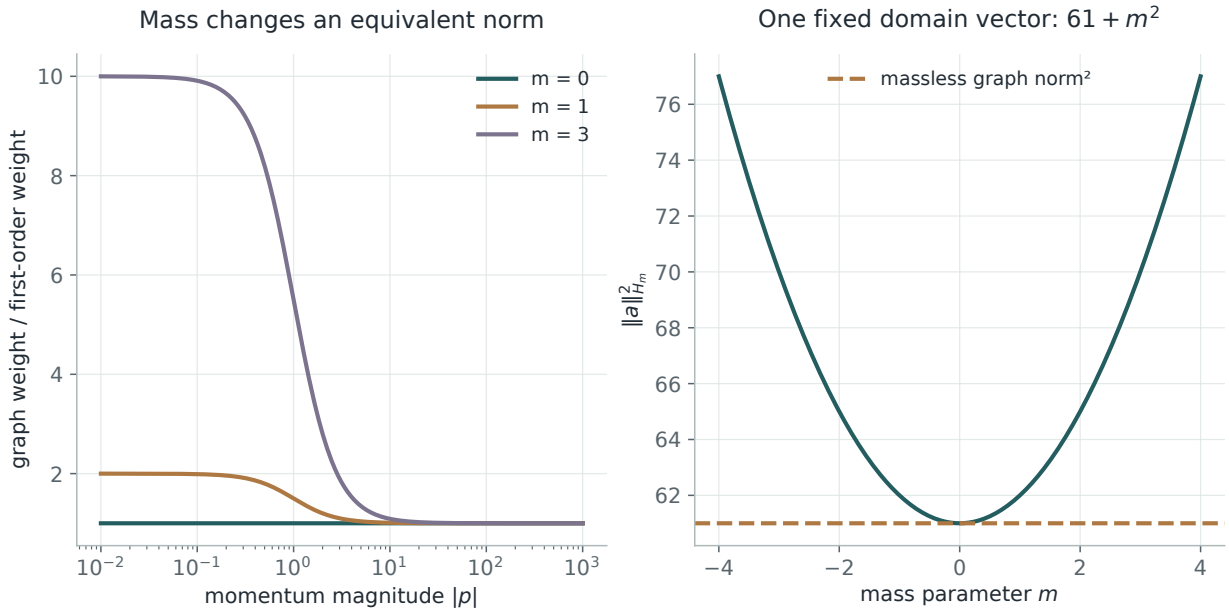
The domain clause is essential. The lower horizontal equivalence is defined everywhere on the Hilbert space, but the vertical arrows need not be. They are unbounded operators with specified dense domains.

Because Fourier transport is unitary, it preserves inner products and transfers self-adjointness and closedness. It also gives the physical form of the weighted domain condition:

$$f \in D(\mathsf{H}_m) \iff \sum_{n\in\mathbb{Z}^3} (1 + |2\pi n|^2) \left\| \hat{f}(n) \right\|^2 < \infty.$$

This is a genuine physical-space operator on torus spinor fields. The square does not by itself identify a previously constructed time-evolution group with the generator of this operator. That relationship is proved in the generator chapter by a full-domain argument, rather than by the diagram alone.

12 Changing mass changes the action, not the domain



Left: the graph weight relative to the first-order Fourier weight for three masses. Right: the exact graph norm square of one fixed normalized domain vector. The physical mass interpretation often takes $m \geq 0$, while the mathematics permits every real m.

For a square-summable coefficient field, the energy identity is

$$\|H_m a\|^2 = m^2 \|a\|^2 + \sum_n |2\pi n|^2 \|a_n\|^2 .$$

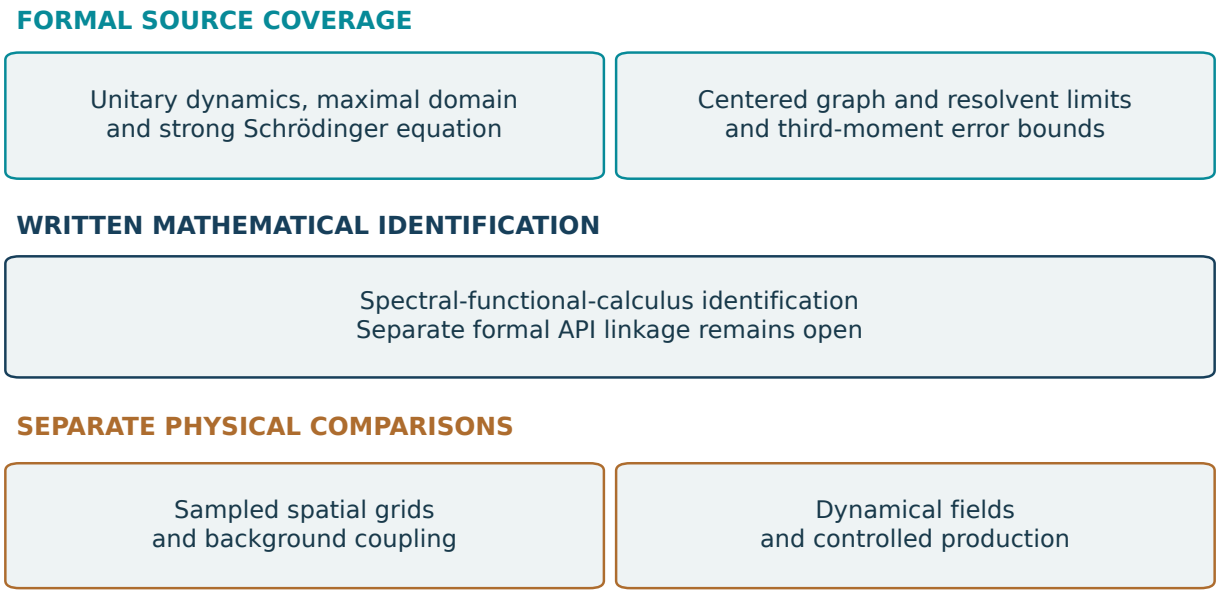
The mass contribution is finite whenever $a \in \ell^2$. Finiteness of the unbounded part therefore depends only on the momentum-weighted sum. Consequently $D(H_{m_1}) = D(H_{m_2})$ for any two real masses.

Equivalent graph norms need not have equal numerical values. For the normalized axis family $a_n = n^{-2}/\sqrt{\zeta(4)}$, the momentum contribution is exactly

$$\frac{4\pi^2\zeta(2)}{\zeta(4)} = 60.$$

The field norm is one, hence $\|a\|^2_{H_m} = 61 + m^2$. A larger mass raises this graph norm while leaving the vector inside the same domain. The graph-weight ratios in the left panel approach one at high momentum because the derivative term dominates.

13 One construction, distinct evidence layers



The checked operator, generator, centered limits and third-order moment estimates, the separate spectral-calculus API linkage, and the comparisons required by an interacting or sampled-grid model. The layers state verification scope rather than physical completeness.

Finite matrices and finite sums supply exact local identities. Completion removes the Fourier cutoff. Fourier transport identifies physical fields, and the adjoint-domain calculation determines the unbounded Hamiltonian. The checked generator theorem then ties the completed group to that operator on the full domain and gives the strong Schrödinger equation. The abstract spectral-functional-calculus identification remains a separate written argument.

For the fixed centered Fourier family, the checked arguments establish graph-domain and strong-resolvent convergence as spacing approaches zero through either sign. A finite third-order spatial Fourier moment gives the quadratic compact-time approximation rate. The technical companion identifies the actual declarations, domain hypotheses, and accepted verification record.

A changing finite spatial lattice needs sampling, interpolation, aliasing control, and compatible geometry. Background coupling changes the operator-domain problem. Dynamical gauge and gravitational fields require additional evolution and constraint arguments. The diagram does not identify those comparisons with the free-field theorem.

14 Reading the plates as one argument

The first three plates fix the representation. Refinement resolves an operator into sectors, normalized Fourier modes identify its spatial channels, and the Clifford matrices fix how each channel acts on a four-component spinor. No convergence claim is needed at this stage. The central facts are exact algebraic identities and exact normalization choices.

The next three plates introduce evolution and approximation. The sine symbol changes the dispersion relation while retaining Hermiticity. Its small-spacing error is quadratic at fixed momentum. Exact modal unitarity then protects the total packet norm even when the local density moves. Stability and consistency serve different purposes: stability prevents growth of an error that is already present, while consistency makes the finite-core error small as the spacing decreases.

The two-tail plate explains how these finite facts reach a full Hilbert space. First approximate a given vector by finite data. Then refine the spacing on that fixed finite approximation. Unitarity controls the discarded part. The argument is powerful precisely because it does not require every square-summable vector to have finite derivative energy.

The remaining operator plates ask a different question: which fields can the continuum Hamiltonian itself act on as an L^2 operator? The weighted-domain criterion answers this exactly. Graph approximation strengthens ordinary approximation by controlling both a field and its image. Fourier transport carries the resulting closed, self-adjoint operator and its domain to physical torus fields. The mass identity shows that this domain reflects spatial regularity rather than a choice of mass.

What the companion can establish. The atlas makes the source architecture inspectable and supplies reproducible examples. The symbolic matrix checks are exact in their finite algebraic setting. The plotted calculations are bounded numerical evaluations with recorded conventions. Infinite-dimensional conclusions depend on the corresponding mathematical arguments and formal source statements.

What remains separate. The figures illustrate the mathematical argument without extending its verification scope. The generator and centered-approximation theorems have written proofs and corresponding formal declarations. The abstract spectral-calculus API linkage remains separate. None of these results supplies the coupled gravitational dynamics or the controlled production history required by the physical matter model.

Gauge Covariance and Geometric Refinement: Visual Atlas

A visual atlas of links, ordered paths, and refinement defects

1 Read the direction before multiplying the links

A link compares vectors that live at different sites. Its direction determines which gauge factor appears on the left and which appears on the right. A plaquette then compares two ways of transporting around a face. Refinement asks a further question: when fine links are composed, do they reproduce the prescribed coarse transporter, and how large is the discrepancy when they do not?

The twelve plates keep these questions separate. The early examples concern exact finite covariance, ordered holonomy and population balance. The later examples concern straight-edge subdivision, selected cubical path comparisons, operator products and regularity bounds. The final plate explains why a geometric refinement needs a matching Hilbert-space normalization.

Two directions, one convention. The older covariant-difference and $SU(2)$ sections use backward links

$$U_{x,\mu} : H_{x+\widehat{\mu}} \longrightarrow H_x, \qquad U'_{x,\mu} = g_x U_{x,\mu} g^{-1}_{x+\widehat{\mu}}.$$

The later refinement construction uses forward links $u = U^{-1}$, with

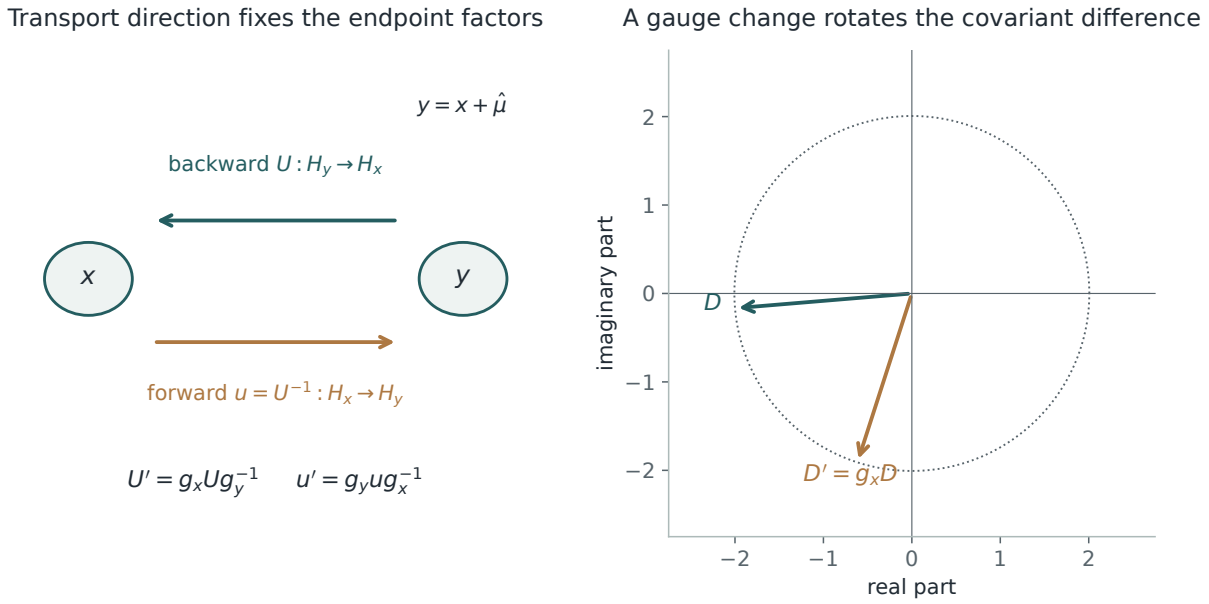
$$u : H_x \longrightarrow H_{x+\widehat{\mu}}, \qquad u' = g_{x+\widehat{\mu}} u g_x^{-1}.$$

States are columns and products act right to left. In particular, forward transport along a first edge u_1 followed by a second edge u_2 is $u_2 u_1$. An Abelian phase commutes with another phase, but its endpoint interpretation still matters.

What a computation establishes. The figure script checks exact rational identities and explicit numerical examples. Every numerical claim has a finite range and a tolerance. The logarithm example uses a small finite family on the principal branch. The operator-product example uses 2×2 matrices on one common space. These illustrations do not replace the general arguments in the monograph or the hypotheses recorded in the Lean companion.

The atlas reports no new Lean compilation. Its figures were made with Python and Matplotlib, and its finite assertions are recorded in a separate computation manifest. The selected formal gauge results establish transport and refinement identities. They do not by themselves supply coupled dynamical gauge-field evolution.

2 A covariant difference belongs to its output site



Left: backward and forward transport have inverse directions and opposite endpoint ordering. Right: one exact rational complex example, plotted in floating coordinates. The covariant difference rotates by the gauge phase at its output site, preserving its magnitude.

For a backward link, define

$$D\psi_x = U_{x,\mu}\psi_{x+\widehat{\mu}} - \psi_x.$$

Both terms lie in H_x. Under the endpoint gauge changes, the intermediate factor cancels:

$$D'\psi'_x = g_x U g^{-1}_{x+\widehat{\mu}} g_{x+\widehat{\mu}}\psi_{x+\widehat{\mu}} - g_x\psi_x = g_x D\psi_x.$$

Dividing by a scalar spacing h gives the usual covariant forward difference without changing this transformation law. In a unitary representation, $\|D'\psi'_x\| = \|D\psi_x\|$.

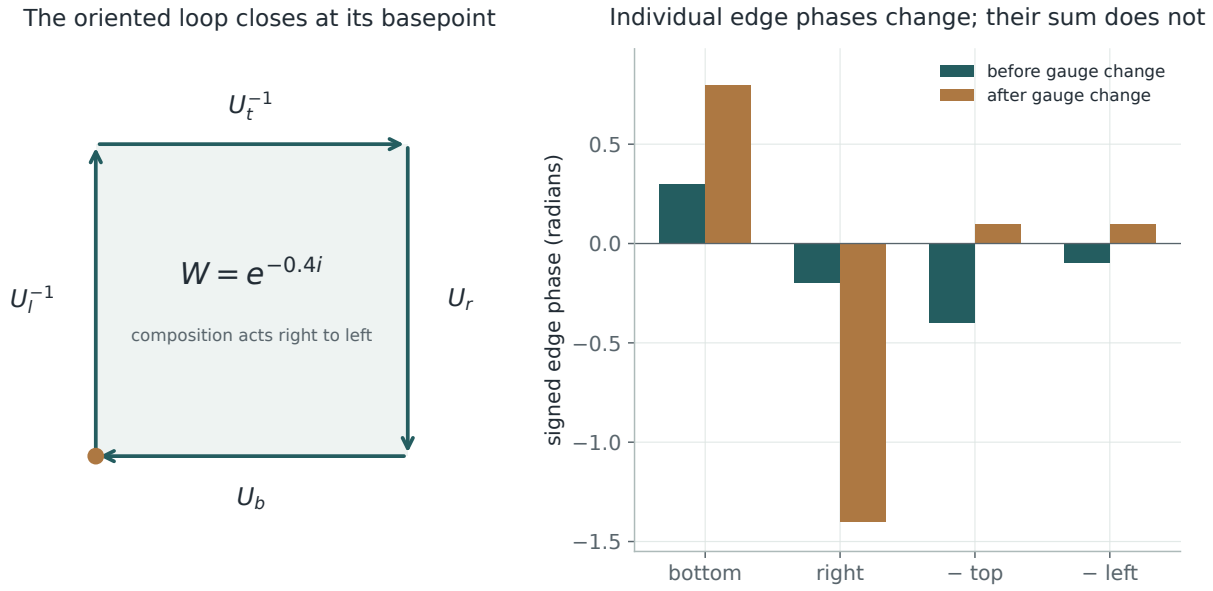
The displayed finite example uses

$$U = \frac{3+4i}{5}, \quad g_x = \frac{5+12i}{13}, \quad g_y = \frac{8-15i}{17}, \quad \psi_x = 1 + \frac{i}{2}, \quad \psi_y = -\frac{1}{3} + i.$$

Each link or gauge factor has unit modulus. The script checks $D' = g_x D$ exactly over the Gaussian rationals before drawing the two vectors. The circle is the locus of vectors with the original magnitude; it illustrates covariance, not a time evolution.

Reversing the link gives $u' = g_y u g_x^{-1}$. Keeping the old endpoint formula after reversing the arrow would assign the transformed vector to the wrong fiber.

3 A closed Abelian loop cancels the endpoint phases



A backward-link plaquette based at the lower-left corner. The arrows follow the composed transport as it acts on a vector. Individual signed edge phases change under a gauge transformation, while their total remains $-2/5$ radians in this finite example.

Write the bottom, right, top and left backward links as U_b, U_r, U_t, U_l. Their based loop is

$$W = U_b U_r U_t^{-1} U_l^{-1}.$$

Applied from right to left, the path goes up the left side, across the top, down the right side and back along the bottom. At every intermediate vertex, the endpoint gauge factor from one link cancels the inverse factor from its neighbor.

For Abelian links $U_e = e^{i\theta_e}$, this reduces to

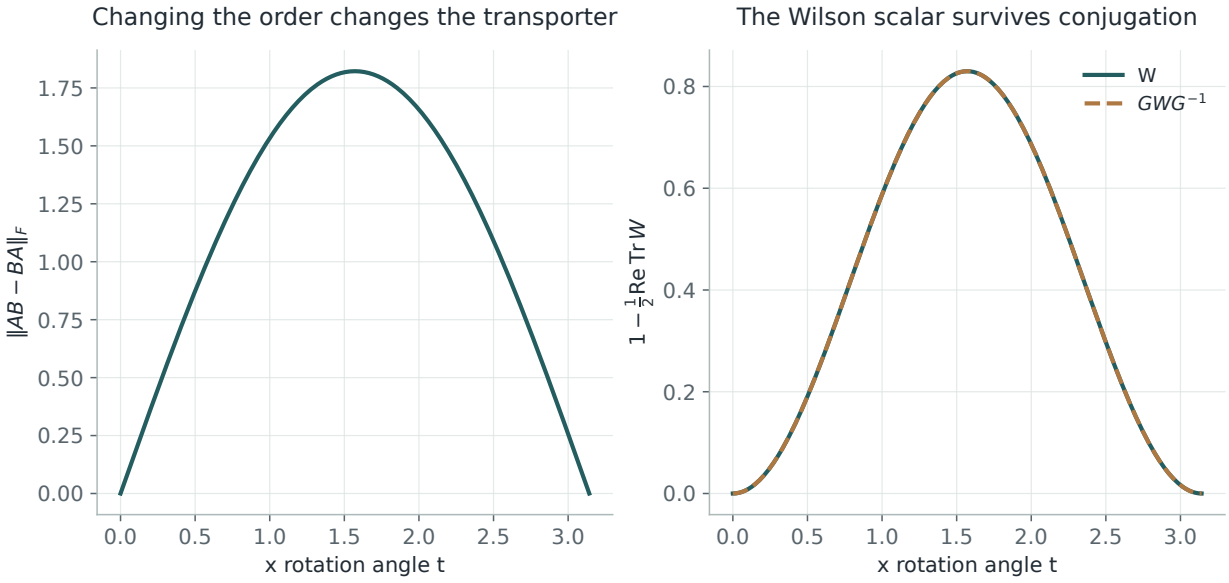
$$W = e^{i\Theta}, \qquad \Theta = \theta_b + \theta_r - \theta_t - \theta_l.$$

The plotted phases are $(3/10, -1/5, 2/5, 1/10)$, giving $\Theta = -2/5$. The gauge angles at the lower-left, lower-right, upper-right and upper-left vertices are $(1/10, -2/5, 4/5, 3/10)$. Their signed changes telescope to zero exactly.

This calculation uses real phase lifts chosen for the example. The observable is the unit-circle value $e^{i\Theta}$, so a general phase angle is defined only modulo 2π. No global principal-angle choice is needed for the exact product identity.

The Wilson scalar $1 - \operatorname{Re} W$ is consequently gauge invariant. Its small-angle expansion can connect a smooth small face to curvature, but that analytic interpretation needs additional smoothness and scaling assumptions.

4 Non-Abelian order matters even when the trace is invariant



A finite family $A(t) = e^{it\sigma_x}$ and $B = e^{0.7i\sigma_y}$. Left: the Frobenius norm of $AB - BA$. Right: the Wilson scalar of $ABA^{-1}B^{-1}$ before and after one fixed gauge conjugation; the curves coincide.

Matrix-valued links do not generally commute. In this example,

$$\|A(t)B - BA(t)\|_{\mathrm{F}} = 2\sqrt{2}\,|\sin t\,\sin(0.7)|.$$

The figure compares this independent formula with the direct matrix calculation at 301 angles. It vanishes at $t = 0$ and $t = \pi$, but is nonzero at intermediate angles. The Frobenius norm here measures the difference between two finite matrices; it is not being identified with a lattice operator norm.

A based loop transforms by conjugation,

$$W' = GWG^{-1},$$

and cyclicity of the trace gives $\mathrm{Tr}\,W' = \mathrm{Tr}\,W$. Thus a gauge-invariant scalar can survive even though the matrix representation of the transporter changes. Gauge invariance of the trace does not allow the internal link factors to be reordered.

The exact benchmark uses rational $SU(2)$ matrices of the form

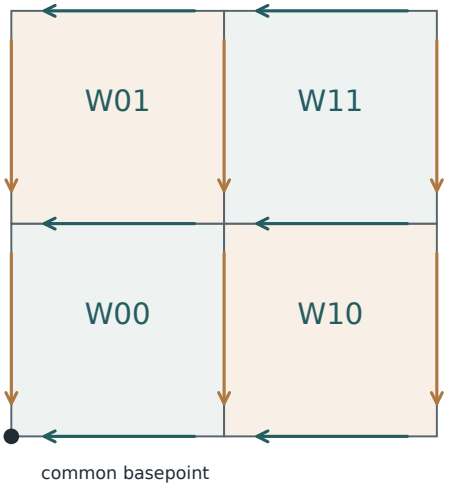
$$Q_a(t) = \frac{(1-t^2)I + 2it\sigma_a}{1+t^2}, \qquad t \in \mathbb{Q}.$$

With the specified rational parameters in the script, unitarity, determinant one, noncommutation and trace invariance are checked algebraically. The smooth curves illustrate the same distinctions over a sampled one-parameter family.

5 Refinement requires a common basepoint and a fixed order

Every interior edge cancels in the right order

Transport to a common basepoint comes first



$$R_j = (X_{0j} W_{1j} X_{0j}^{-1}) W_{0j}$$

$$W_{\text{boundary}} = R_0 Y_{00} R_1 Y_{00}^{-1}$$

Ordered reconstruction

exact residual: 0

Naive product of four based plaquettes

Frobenius error: 2.7587

A 2×2 array of noncommuting rational $SU(2)$ plaquettes. The source refinement formula reproduces the eight-link coarse boundary exactly. A naive product of the four plaquettes, without their required basepoint transports, differs in this same example.

Let X_{ij} and Y_{ij} denote backward horizontal and vertical fine links. Define

$$W_{ij} = X_{ij} Y_{i+1,j} X_{i,j+1}^{-1} Y_{ij}^{-1}.$$

The right plaquette in a row must first be moved to the left basepoint. The correct row loop and coarse loop are

$$R_j = (X_{0j} W_{1j} X_{0j}^{-1}) W_{0j}, \qquad W^{(2)} = R_0 Y_{00} R_1 Y_{00}^{-1}.$$

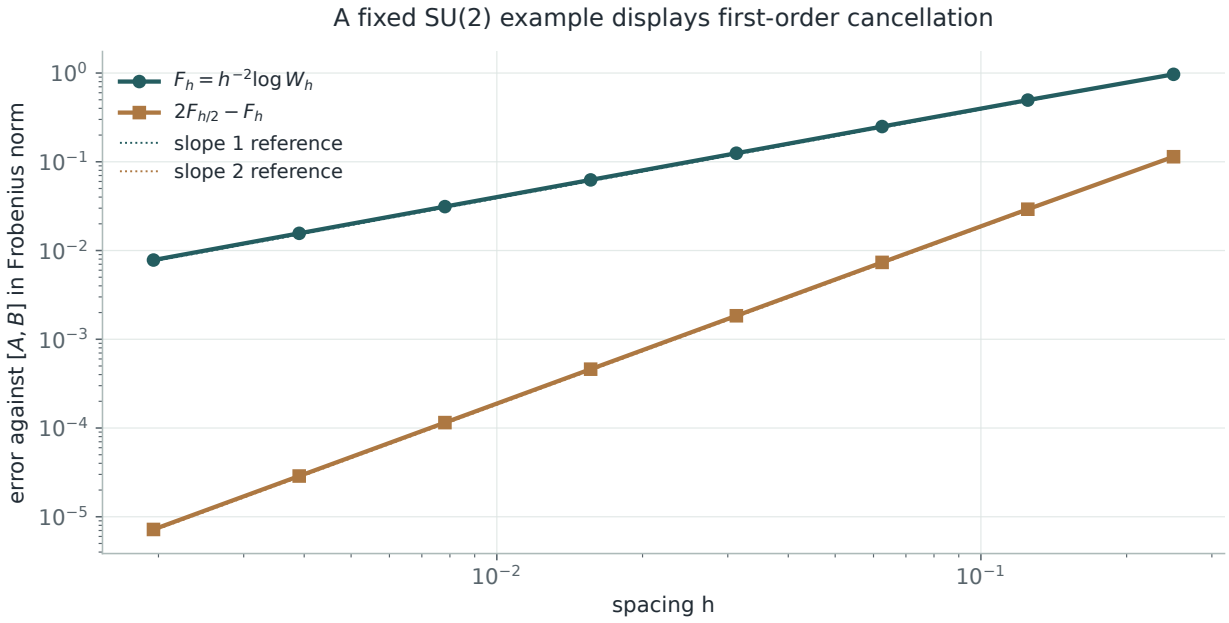
Expanding these factors cancels each interior link with its adjacent inverse. The surviving boundary is

$$X_{00} X_{10} Y_{20} Y_{21} X_{12}^{-1} X_{02}^{-1} Y_{01}^{-1} Y_{00}^{-1}.$$

No commutation of distinct matrices is used. Moving the factors merely because the cells look interchangeable would destroy the cancellation.

All twelve edge matrices in this plate are specified by rational Cayley parameters in the script. Exact Gaussian rational arithmetic gives zero for the difference between the boundary product and the transported row expression. The naive product $W_{00} W_{10} W_{01} W_{11}$ is different. Its plotted residual is a counterexample to that naive finite prescription, not to the correct refinement identity.

6 A first-order curvature error can be cancelled



The explicit matrices $A = i\sigma_x$ and $B = i\sigma_y$, with $h = 2^{-2}, \ldots, 2^{-9}$. The direct logarithmic curvature approaches $[A, B]$ with an observed first-order slope, while the displayed Richardson combination has an observed second-order slope. These finite slopes do not prove the asymptotic expansion.

Consider

$$W_h = e^{hA}e^{hB}e^{-hA}e^{-hB}, \qquad F_h = h^{-2}\log W_h.$$

On a common analytic logarithm branch near the identity, the expansion has the form

$$\log W_h = h^2[A, B] + h^3C_3 + O(h^4).$$

Therefore $F_h = [A, B] + hC_3 + O(h^2)$. Evaluating the same family at half spacing and subtracting gives

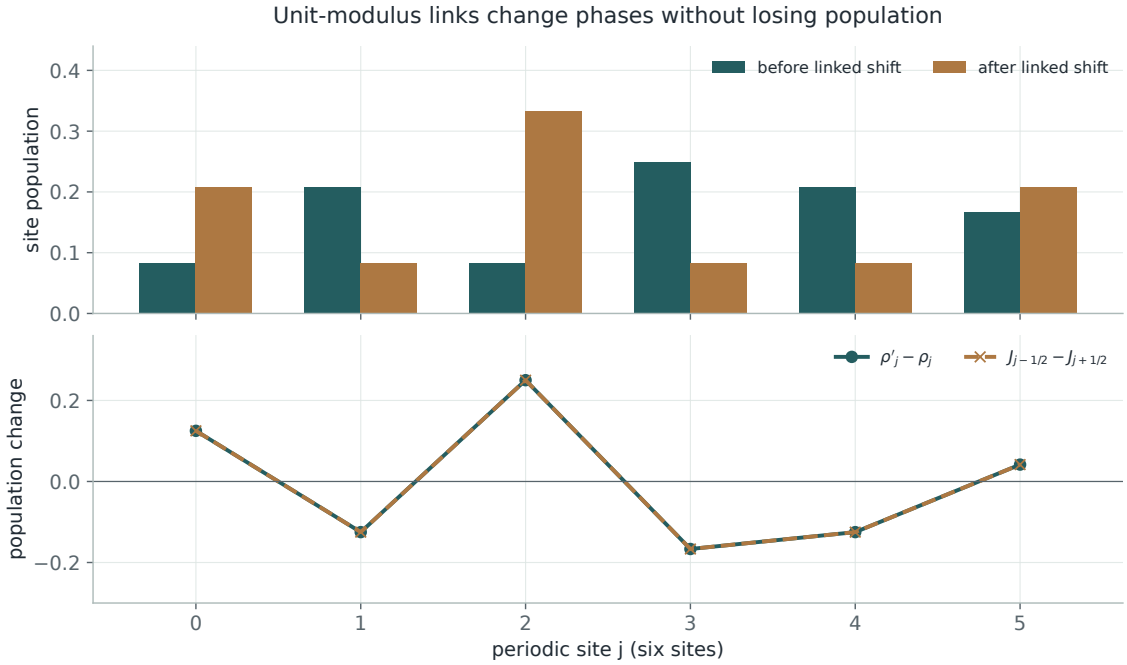
$$2F_{h/2} - F_h = [A, B] + O(h^2).$$

The cancellation follows from the expansion, not from the appearance of the plotted lines.

For this example, $[A, B] = -2i\sigma_z$. The script computes the principal matrix logarithm and records the largest absolute eigenangle for every displayed matrix and every half-spacing matrix used by the extrapolation. This verifies that the finite calculation stays near the identity; it does not establish a branch condition for every possible matrix family.

At much smaller spacings, division by h^2 can amplify floating error. The present fixed range stops before that effect controls the plot. No larger range is claimed, and the numerical logarithm is not interval certified.

7 A linked shift has an exact local population balance



One linked shift on six periodic sites, starting from explicitly specified post-coin amplitudes. The lower curves compare the observed site-population change with incoming current minus outgoing current. Their numerical residual is checked against the recorded tolerance.

Let η_j^+ and η_j^- be the two post-coin components. For backward unit-modulus links U_j, the shift is

$$\xi_j^+ = U_{j-1}^{-1}\eta_{j-1}^+, \qquad \xi_j^- = U_j\eta_{j+1}^-.$$

Its current through the edge between j and $j+1$ is

$$J_{j+1/2} = |\eta_j^+|^2 - |\eta_{j+1}^-|^2.$$

Writing $\rho_j = |\eta_j^+|^2 + |\eta_j^-|^2$ and defining ρ'_j from ξ, unit modulus immediately gives

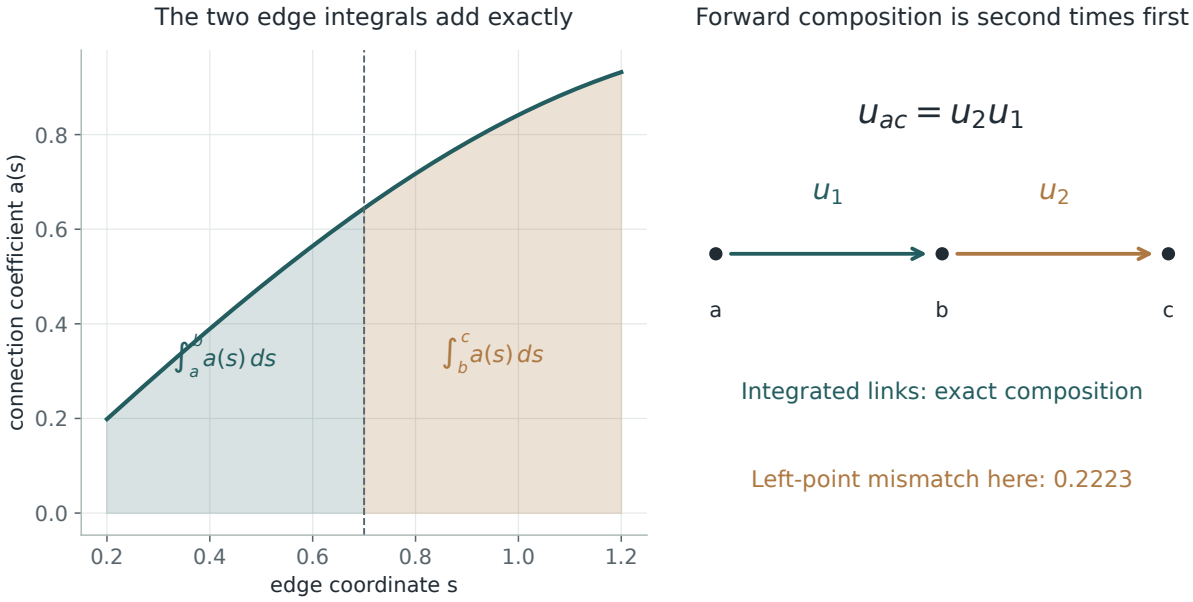
$$\rho'_j - \rho_j = J_{j-1/2} - J_{j+1/2}.$$

Summing over the periodic lattice cancels all currents, so the total population is preserved.

The example uses post-coin vectors proportional to $(1, 2i, -1, 1+i, -2i, 2)$ and $(i, 1, -i, 2, -1, 0)$, both divided by $\sqrt{24}$. Their combined squared norm is one. The six link phases are listed in the computation manifest. This plate describes one finite step; it makes no assertion about a separately chosen coin sequence or a coupled gauge-field update.

The balance law also explains why the current is signed. A negative edge current represents net transport in the reverse direction, rather than a negative probability.

8 Integrated edge links compose under subdivision



An Abelian connection coefficient $a(s) = \sin s$ on the interval $[0.2, 1.2]$, split at 0.7. The two exact edge integrals add to the coarse integral. Replacing them by independent left-point samples produces a nonzero mismatch in this same example.

For the forward convention in this plate, set

$$u_{ab} = \exp\left(i \int_a^b a(s)\, ds \right).$$

Integral additivity gives

$$u_{ac} = u_{bc} u_{ab}.$$

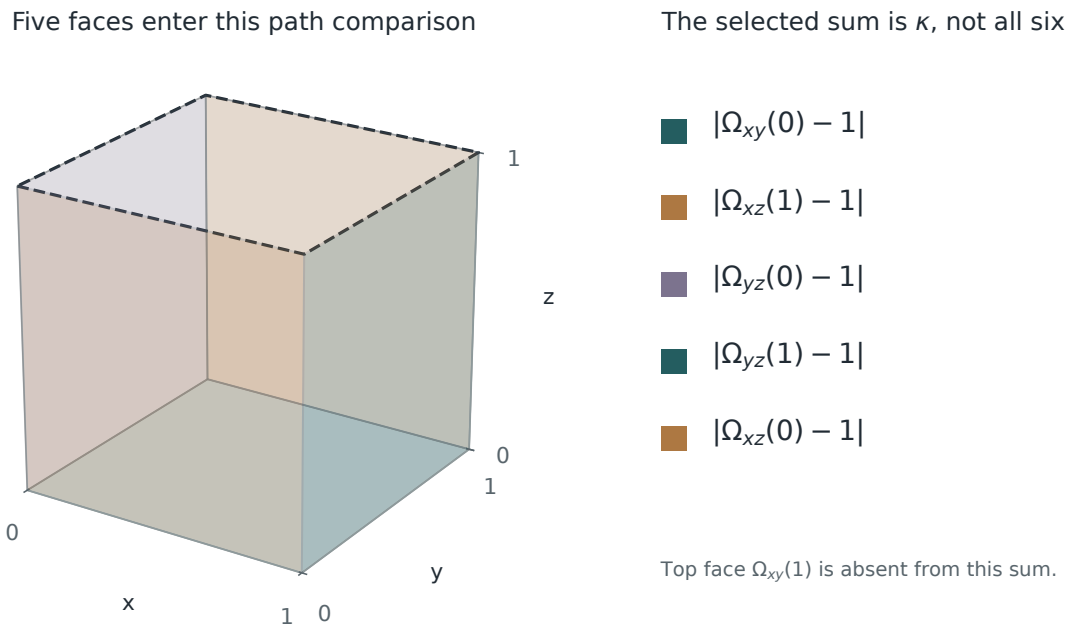
The order is second edge times first edge. The phases commute in this Abelian example, but the direction and endpoints are retained so that the formula remains compatible with the general transport convention.

For $a(s) = \sin s$, the integral is $\cos a - \cos b$. The script compares the coarse phase with the product of the two fine phases using this analytic expression. Their small floating residual measures evaluation error in an exactly composing example.

By contrast, a left-point prescription sets $\tilde{u}_{ab} = e^{i(b-a)\sin a}$. Its coarse exponent and the sum of its two fine exponents need not agree. This does not make left-point sampling unusable; it means exact subdivision is an additional property that this prescription does not automatically have.

Similarly, a centered second-order continuum formula needs suitable link sampling, such as midpoint data or exact edge integrals under the stated regularity assumptions. A spatially varying left-point coefficient alone does not justify a second-order claim.

9 The cubical comparison selects five face defects



The five faces used by the fixed cubical path comparison. The dashed top face $xy(1)$ is not a term in this budget. This plate is a combinatorial diagram, with an inventory check; it is not an operator-norm computation.

In the inherited cubical comparison, the curvature contribution is

$$\kappa = |\Omega_{xy}(0) - 1| + |\Omega_{xz}(1) - 1| + |\Omega_{yz}(0) - 1| \\ + |\Omega_{yz}(1) - 1| + |\Omega_{xz}(0) - 1|.$$

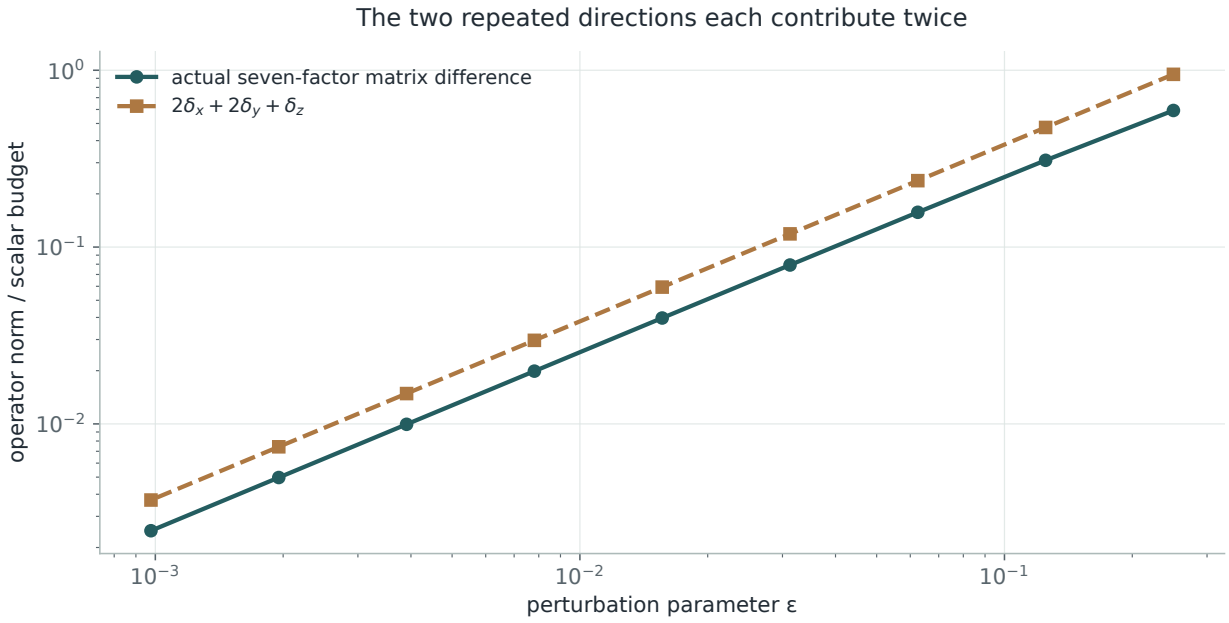
Each term measures the unit-circle phase defect of a selected face. The arguments 0 and 1 specify the fixed value of the coordinate transverse to that face. Thus $xy(0)$ is the bottom face and $xy(1)$ is the top face.

The chosen comparison decomposes its change of paths through five face swaps. The budget records those swaps. It is not the sum over every face of the cube, and it is not a claim that the omitted face is physically irrelevant. A different path comparison can have a different useful decomposition.

When the selected face phases are all one, this budget vanishes. With nontrivial curvature it bounds the corresponding path discrepancy under the stated operator hypotheses. The Lean companion identifies the exact construction and the factor ordering on which the bound depends.

The scalar inequality $|e^{i\theta} - 1| \le |\theta|$ later converts regularity estimates for real face phases into bounds for these unit-circle defects. Keeping the phase defect separate from its real-angle upper bound avoids imposing an unnecessary global logarithm.

10 Repeated factors contribute repeated error terms



A common-space 2×2 unitary surrogate for the product pattern $MXYZYXM$. The actual spectral-norm difference stays below the repeated-factor budget at nine specified perturbation sizes. The matrices illustrate telescoping; they are not the formal Clifford lattice walk.

For contractive factors on a common normed space, subtracting two products one factor at a time gives a telescoping sum. Each term contains one local factor difference surrounded by contractions, so its norm is no larger than that local difference.

Here the two products have the pattern

$$W = MXYZYXM, \qquad \widetilde{W} = M\widetilde{X}\widetilde{Y}\widetilde{Z}\widetilde{Y}\widetilde{X}M.$$

The mass factors agree. If $\delta_x = \left\|\widetilde{X} - X\right\|$ and similarly for the other directions, then

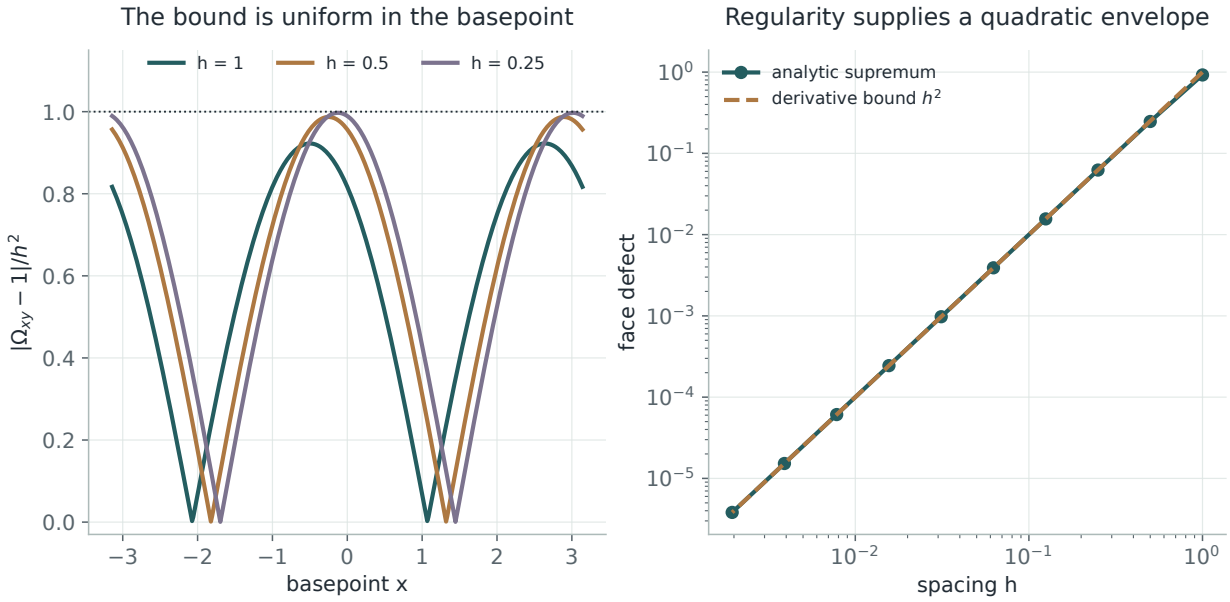
$$\left\|\widetilde{W} - W\right\| \leq 2\delta_x + 2\delta_y + \delta_z.$$

The coefficients two count actual occurrences in the product. They cannot be dropped just because only three direction names appear.

The script specifies all seven finite matrices through Pauli exponentials and measures the actual operator differences. It checks unitarity and the bound for $\varepsilon = 2^{-2}, \ldots, 2^{-10}$. These δ values are matrix-factor differences; the formal lattice theorem has a separate argument connecting its local transport defects to edge-ratio suprema.

For different fine and coarse spaces, an embedding or comparison map must first be included in the local defects. The present common-space example cannot supply that missing identification by itself.

11 A regular connection supplies a uniform quadratic bound



For the prescribed background $A_y = \sin x$ and $A_x = A_z = 0$, the face phase is $h[\sin(x+h) - \sin x]$. Left: normalized defects at three spacings. Right: the analytic supremum and the derivative envelope h^2 for $h = 1, 1/2, \ldots, 2^{-9}$.

Exact edge integration gives

$$\Omega_{xy}(x,h) = \exp(ih[\sin(x+h) - \sin x]).$$

Since $|\sin(x+h) - \sin x| \le h$ for $h \ge 0$,

$$|\Omega_{xy}(x,h) - 1| \le h^2.$$

This is uniform in the basepoint. The estimate follows from the derivative bound $|\cos x| \le 1$ and $|e^{it} - 1| \le |t|$, rather than from checking a sufficiently dense grid.

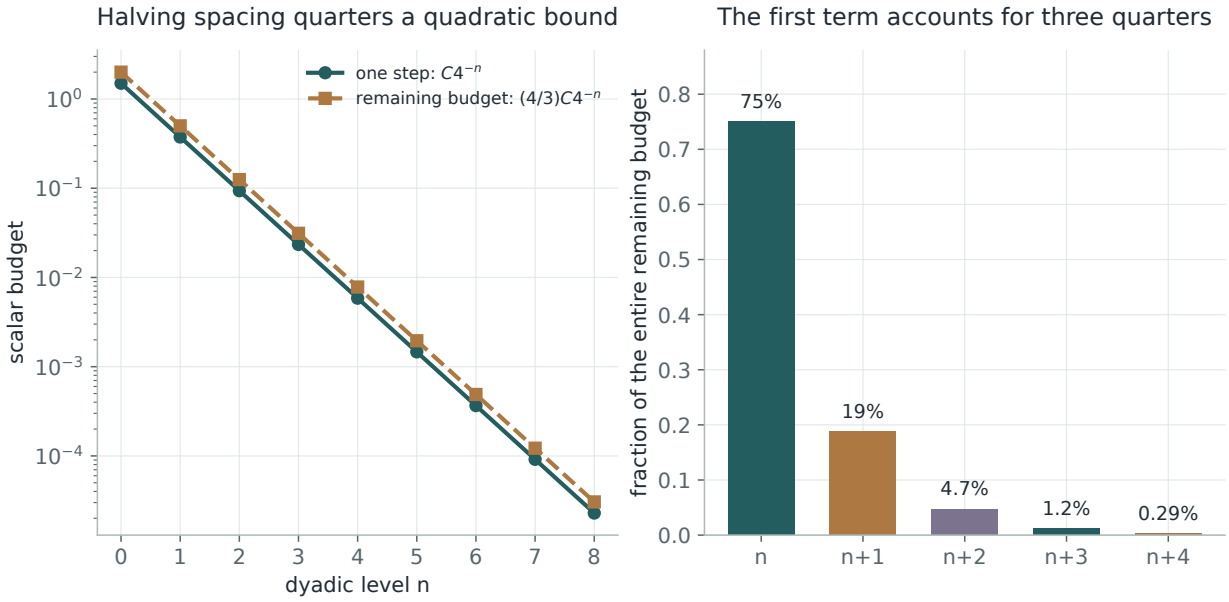
The trigonometric identity

$$\sin(x+h) - \sin x = 2\sin(h/2)\cos(x + h/2)$$

also gives the exact supremum $2\sin(h\sin(h/2))$ when $0 < h \le 1$. The script compares 401 sampled basepoints with that analytic supremum and with the simpler h^2 envelope. This supplies an independent check of the plotted implementation.

All other coordinate face types are flat for this prescribed coefficient field. Exact integrated links also have zero straight-edge subdivision defect. The example therefore isolates a curvature contribution without introducing a coupled gauge evolution. A background regularity estimate does not prove that an evolving gauge field retains those bounds.

12 Quadratic dyadic errors have a summable remaining budget



The explicit scalar model $b_n = C4^{-n}$ with $C = 3/2$. Left: one-step and remaining geometric budgets. Right: the first five contributions as fractions of the entire remaining budget. The infinite-tail formula is derived analytically; the script checks exact rational finite decompositions.

If the spacing is $h_n = a_0\, 2^{-n}$, a uniform bound proportional to h_n^2 becomes

$$b_n = C4^{-n}.$$

The scalar geometric series has the tail

$$\sum_{k=n}^{\infty} b_k = \frac{4}{3} C4^{-n}.$$

Its first term is three quarters of the whole tail. The following term contributes $3/16$, the next $3/64$, and so on. The displayed percentages are evaluations of these exact fractions.

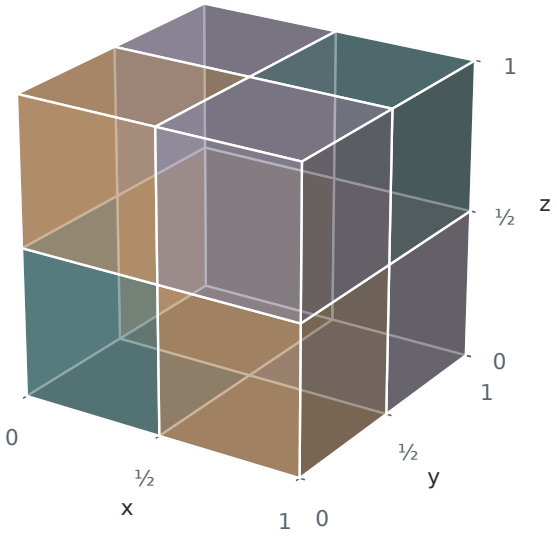
For every displayed starting level n, the script checks that the finite sum through level eight plus the stated level-nine tail equals the full level-n tail. These are exact rational identities. They validate the formula's implementation, while the universal tail identity comes from the geometric-series argument.

To convert such a scalar tail into a Cauchy estimate for states or operators, all consecutive differences must be represented in a compatible common space and satisfy the corresponding norm bounds. The scalar sum alone does not construct the embeddings, identify the limiting physical operator, or establish an equation for a coupled continuum theory.

The distinction is useful: a summable local budget is a strong ingredient, but the compatibility of the objects being compared remains part of the proof.

13 Geometric subdivision needs a compatible normalization

Eight child cubes fill one parent cube



Cell volume and amplitude normalization agree

Volume-weighted cell values

$$\sum_{j=1}^{8} \tfrac{1}{8}\|v\|^2 = \|v\|^2$$

Unweighted coefficient coordinates

$$\sum_{j=1}^{8} \|v/\sqrt{8}\|^2 = \|v\|^2$$

The normalization belongs to the chosen representation.

A single unit cube divided into eight equal children. Repeating the same cell value preserves a volume-weighted norm. In unweighted coefficient coordinates, the same embedding divides each child amplitude by $\sqrt{8}$. This is one exact finite isometry example.

Suppose a parent cube has volume one and each child has volume $1/8$. If the parent carries a vector value v, repeating that value on every child gives

$$\sum_{j=1}^{8} \frac{1}{8} \|v\|^2 = \|v\|^2 .$$

Here the geometric volume is part of the norm. If the child coordinates instead use the ordinary unweighted Euclidean norm, the corresponding isometric embedding is

$$v \longmapsto \left(\frac{v}{\sqrt{8}}, \dots, \frac{v}{\sqrt{8}} \right) .$$

The two formulas express the same normalization in different coordinates. Repeating v in an unweighted eight-child space would multiply its squared norm by eight.

The exact example uses $v = (2, -1 + i)^T$, with $\|v\|^2 = 6$. The script checks both child representations and the parent value with rational and radical arithmetic. The cube drawing shows only this one subdivision; it is not a proof about arbitrary partitions.

Finite-cell interpolation, weighted reconstruction and concrete Euclidean box realization are separate steps. A transported finite-resolution Euclidean subspace is not automatically the ordinary dyadic-box subspace. That identification requires its own construction and proof.

14 Reproduce the plates and read the evidence

From the extracted source root, run

```
python tools/make_gauge_figures.py
```

The script writes twelve PDF plates and twelve PNG previews to `source/gauge_figures/`. The successful computation record is `evidence/gauge/figure_checks.json`; the append-only run list is `evidence/gauge/figure_run_history.jsonl`. It records the source hash, before-and-after input identity, actual software dependencies, deterministic input families, numerical tolerances, artifact checksums and elapsed time.

Exact finite checks. The Gaussian rational calculations cover endpoint covariance, signed phase cancellation, a noncommuting $SU(2)$ benchmark, conjugation invariance and the ordered 2×2 refinement identity. Rational or radical checks also cover the scalar dyadic decompositions and eight-child normalization. These calculations rely on SymPy and are not proof-assistant compilation receipts.

Numerical checks. The floating calculations cover the plotted $SU(2)$ family, principal-log curvature values, one six-site linked walk, integrated edge phases, one common-space operator-product family and the regular sinusoidal connection. Their scopes and tolerances are written beside each check. Matrix logarithms, spectral norms and trigonometric evaluations are not interval certified.

The numerical provenance record retains failed and corrected checks as well as successful outputs. A successful bounded calculation is evidence for its specified input and tolerance, not a universal proof.

Interpretation. The figures make the finite mechanisms inspectable: endpoint cancellation, ordered interior-link cancellation, local population balance, integral additivity, repeated-factor telescoping and norm-preserving subdivision. The monograph supplies the mathematical arguments and their restrictions. The Lean companion identifies the inherited formal evidence. None of the three evidence forms silently upgrades the others.

Constrained Gravity and Strong Fields: Visual Atlas

A visual atlas of local balance, regular reduction, and horizons

1 Read the constraint and its domain together

A finite identity can hold exactly while its physical interpretation requires additional assumptions. A constrained step can preserve its invariants algebraically while a numerical implementation retains a nonzero solve residual. A horizon can be regular in one coordinate chart while a different regularity boundary still limits that chart. The twelve plates make these distinctions visible.

The first five concern local scalar work, conformal holonomy, a polarized Gowdy mode, the centered linearized ADM symbol, and a finite nonlinear bracket obstruction. The next three concern regular reduction and a discrete Dirac-gradient step. The final four concern an explicitly chosen quasienergy model, the spherical electrovac family, scalar mass flux, and inherited first-horizon data.

What was computed for the documented audit. All figures were generated with Python and Matplotlib. Exact checks use rational matrices or symbolic polynomial expressions. The sampled families use double precision. A fresh constrained-particle computation illustrates the implicit step on a six-dimensional mechanical example. Its energy and constraint errors are numerical results with stated tolerances; it is not a gravitational evolution.

The last plate redraws a previously archived scalar-collapse event table. Its data are pinned by a checksum and their mass identities are checked, but the collapse evolution, radial projection, and event extraction are not rerun. An inherited event is therefore labeled as inherited throughout.

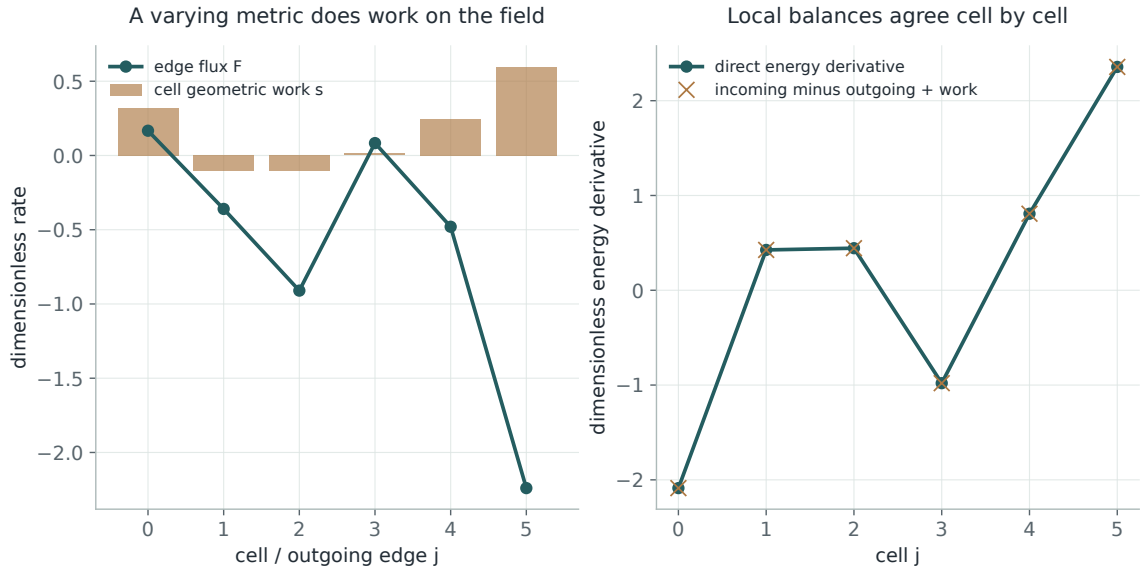
Conventions. The conformal cell has positive orientation $dt \wedge dx$. Periodic centered differences use $D_h f_j = (f_{j+1} - f_{j-1})/(2h)$. Canonical particle coordinates are ordered $(q_1, q_2, q_3, p_1, p_2, p_3)$, with

$$J = \begin{pmatrix} 0 & I \\ -I & 0 \end{pmatrix}.$$

Spherical strong-field plates use the ingoing Painlevé–Gullstrand choice $v \geq 0$. Their units are $G = c = 1$, with $4\pi\epsilon_0 = 1$ for electrovac. For the scalar sector, $\varphi = \sqrt{4\pi}\,\phi_{\text{phys}}$.

Formal boundary. The gravity chapters are supported by written derivations and finite computations at their stated scope. The retained formal free-Dirac construction does not certify the gravitational model. The verification companion records the distinction between these kinds of evidence.

2 A changing metric contributes geometric work



A six-cell periodic scalar system with rational field, velocity, and coefficient data. Left: interface flux and cell work. Right: the directly differentiated cell energy agrees with incoming flux minus outgoing flux plus geometric work. The equality is checked exactly before plotting.

Let $w_j = a_j/N_j$, $c_{j+1/2} = N_{j+1/2}/a_{j+1/2}$ and $v_j = N_j a_j m^2$. The finite action gives

$$w_j\ddot{\varphi}_j + \dot{w}_j\dot{\varphi}_j = \frac{c_{j+1/2}\Delta_j - c_{j-1/2}\Delta_{j-1}}{h^2} - v_j\varphi_j, \qquad \Delta_j = \varphi_{j+1} - \varphi_j.$$

With half of each neighboring edge energy assigned to its cell,

$$e_j = \frac{h}{2}(w_j\dot{\varphi}_j^2 + v_j\varphi_j^2) + \frac{c_{j+1/2}\Delta_j^2 + c_{j-1/2}\Delta_{j-1}^2}{4h},$$

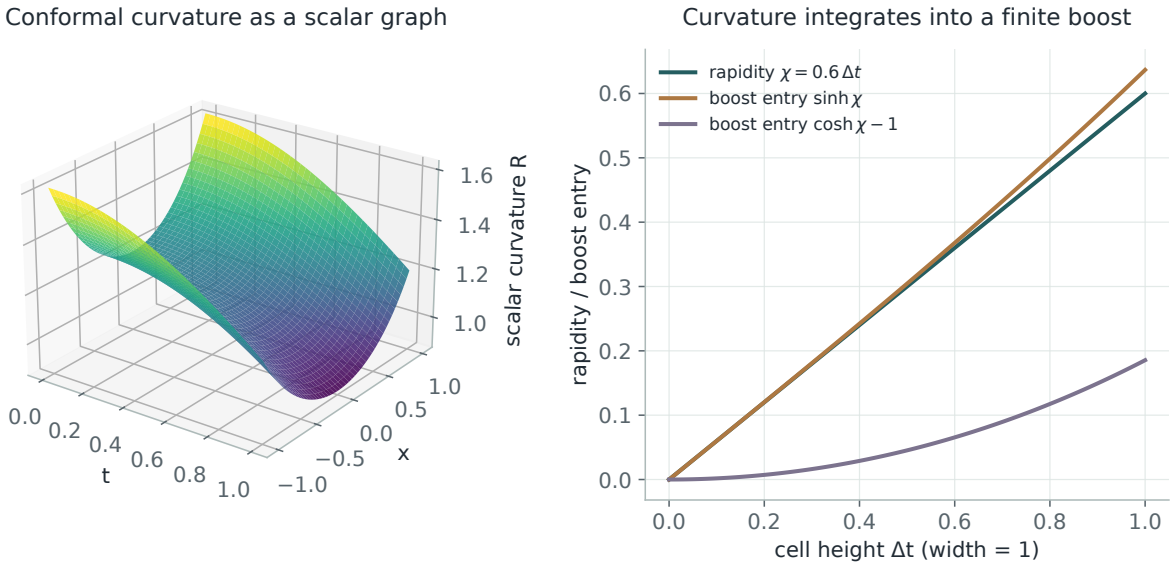
the local law is

$$\dot{e}_j + F_{j+1/2} - F_{j-1/2} = s_j, \qquad F_{j+1/2} = -\frac{c_{j+1/2}\Delta_j}{2h}(\dot{\varphi}_j + \dot{\varphi}_{j+1}).$$

The work term contains $-h\dot{w}_j\dot{\varphi}_j^2/2$, $h\dot{v}_j\varphi_j^2/2$, and the two coefficient derivatives of the edge energies. Thus a time-dependent prescribed geometry need not conserve the scalar-field energy alone. On a periodic chain, the fluxes telescope and $d\sum_j e_j/dt = \sum_j s_j$.

The example specifies $h = 1/2$, $w_j = 1 + j/10$, $c_{j+1/2} = 1 + (j \bmod 3)/5$, and $v_j = 1/3$. The script records every remaining rational input. Its six local residuals and the summed residual are exactly zero. This is an instantaneous identity for a finite scalar system on prescribed metric coefficients. It does not include an independently evolved gravitational energy.

3 Conformal curvature integrates into Lorentzian holonomy



Left: a genuine three-dimensional graph of the scalar curvature over (t, x), not a spacetime embedding. Right: the integrated rapidity and two entries of the finite Lorentz boost for cells of width one and varying height.

For

$$ds^2 = e^{2\phi}(-dt^2 + dx^2), \qquad \omega = (\phi_x\, dt + \phi_t\, dx)K, \qquad K = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix},$$

all connection values are proportional to the same generator. Consequently,

$$\Omega = (\phi_{tt} - \phi_{xx})K\, dt \wedge dx, \qquad R = 2e^{-2\phi}(\phi_{tt} - \phi_{xx}).$$

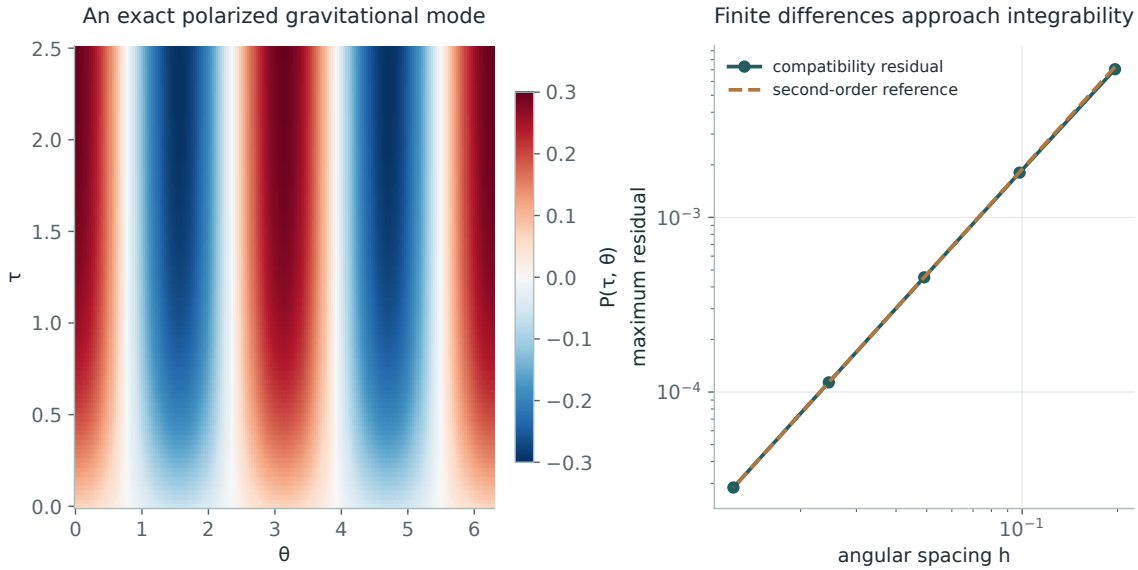
This plate uses signed plus-edge links $U_e = \exp(+\int_e \omega)$. Their product around a positively oriented rectangle gives

$$W_\square = e^{\chi K}, \qquad \chi = \int_\square (\phi_{tt} - \phi_{xx})\, dt\, dx = \int_\square \tfrac{1}{2} e^{2\phi} R\, dt\, dx.$$

This example chooses $\phi = 3(t^2 - x^2)/20$. Then $\phi_{tt} - \phi_{xx} = 3/5$ and $R = (6/5)e^{-2\phi}$. The cell $[0, 1] \times [-1/2, 1/2]$ has $\chi = 3/5$. Independent symbolic boundary and area integrals agree exactly.

Because $K^2 = I$, the signed loop is $e^{\chi K} = I \cosh\chi + K \sinh\chi$. With the Cartan convention $D = d + \omega$, ordinary parallel transport is $\exp(-\int_e \omega)$ and the loop has the opposite rapidity. Both are Lorentz boosts, rather than unitary matrices in the positive Euclidean norm. The surface displays varying curvature while $e^{2\phi}R/2$ is constant. The identity uses this conformal reduction and commuting connection generator.

4 An exact Gowdy mode tests constraint integrability



Left: the exact polarized mode with $A = 0.3$ and $n = 2$. Right: at $\tau = 0.7$, the maximum centered-difference compatibility residual for $N = 32, 64, 128, 256, 512$. The reference line has slope two; the measured last-pair order is recorded in the computation manifest.

In the polarized sector $Q = 0$, the reduced field equation is

$$P_{\tau\tau} - e^{-2\tau} P_{\theta\theta} = 0.$$

Writing $x = ne^{-\tau}$ gives the exact periodic family

$$P = AJ_0(x)\cos(n\theta), \qquad P_\tau = AxJ_1(x)\cos(n\theta).$$

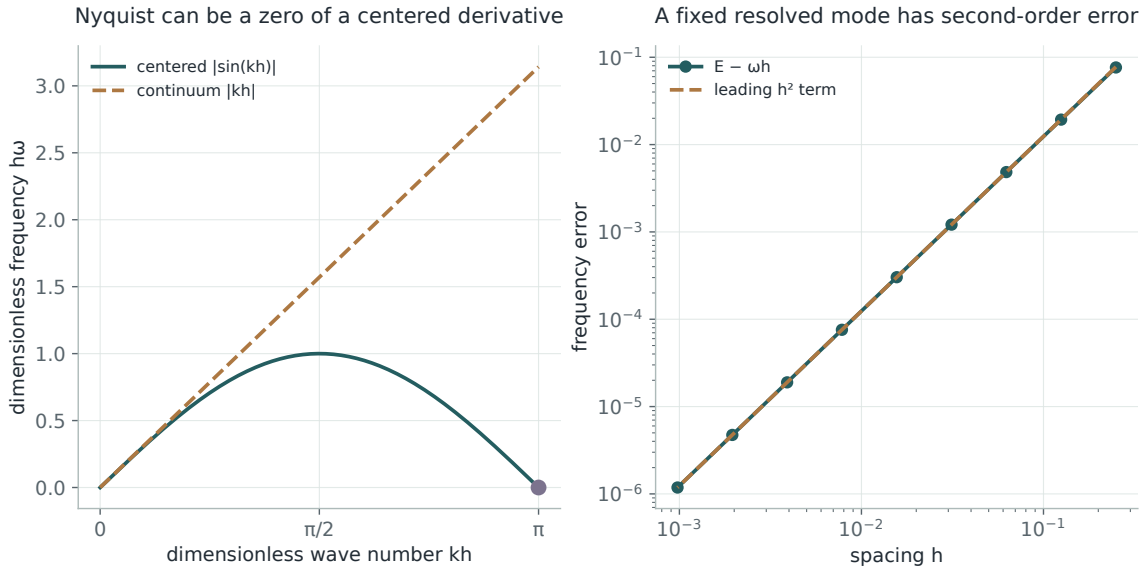
The Bessel equation supplies $P_{\tau\tau} = -x^2 P$, so the wave equation holds identically. The reduced constraint densities become

$$\lambda_\tau = A^2 x^2 \left[J_1(x)^2 \cos^2(n\theta) + J_0(x)^2 \sin^2(n\theta)\right],$$
$$\lambda_\theta = -A^2 n x J_0(x) J_1(x) \sin(2n\theta).$$

Their exact identity is $\partial_\tau \lambda_\theta = \partial_\theta \lambda_\tau$. The script evaluates the left side analytically and replaces only the right side by a centered difference. The resulting nonzero residual belongs to the chosen finite approximation. The periodic momentum integral vanishes because n is an integer and the density is proportional to $\sin(2n\theta)$.

The plot is an analytic polarized solution, with a bounded consistency study. It does not show a newly integrated unpolarized trajectory, a global nonlinear Einstein solution, or a claim about all grid resolutions. The distinction matters because global endpoint projection and local continuum compatibility are separate conditions in the semidiscrete unpolarized system.

5 The centered ADM symbol includes a Nyquist null



Left: the one-dimensional centered-derivative symbol and its continuum comparison. The Nyquist point is a zero of the centered symbol. Right: the fixed three-dimensional mode $k = (1, 2, 0.5)$ approaches its continuum frequency under dyadic spacing refinement.

On a periodic Fourier mode, the centered spatial derivative acts by multiplication with $i\tilde{k}_i$, where

$$\tilde{k}_i = \frac{\sin(k_i h)}{h}.$$

The transverse-traceless sector of the linearized projector–ADM equations satisfies $\ddot{h}_{ij} = \delta^2 h_{ij}$ and hence

$$\omega_h^2 = \sum_i \tilde{k}_i^2.$$

For a fixed nonzero resolved wave vector,

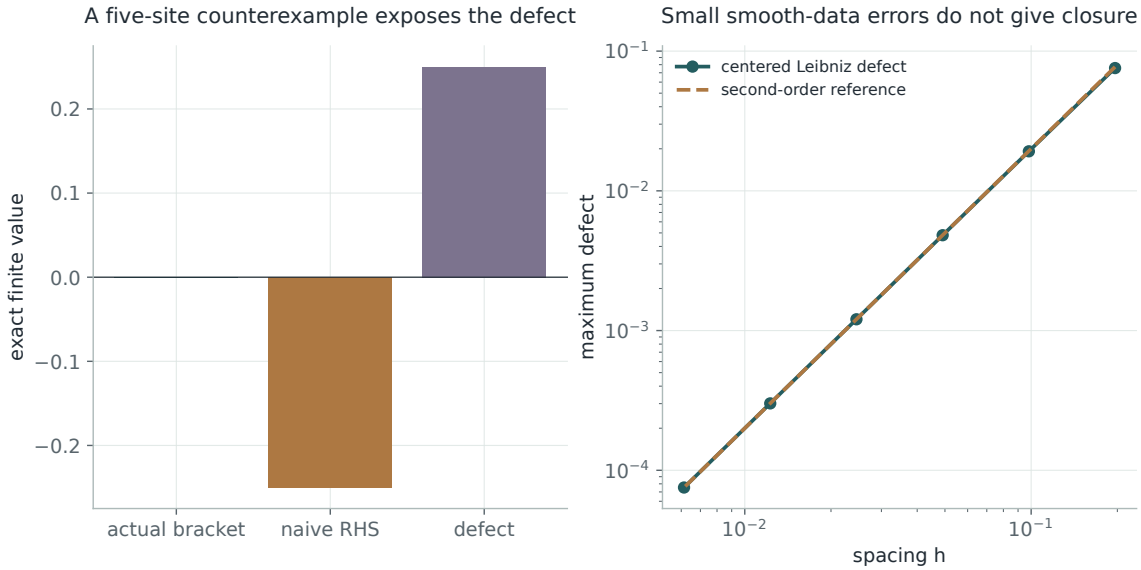
$$\omega_h = |k| - \frac{h^2}{6|k|} \sum_i k_i^4 + O(h^4).$$

The numerical panel compares the exact finite symbol with this leading correction for $h = 2^{-2}, \ldots, 2^{-10}$. Its observed order concerns this fixed vector and finite range.

At $kh = \pi$, the sampled checkerboard $(-1)^j$ has centered derivative zero. In several dimensions, a vector whose components are individually zero or Nyquist can similarly have $\tilde{k} = 0$. A transverse projector that divides by $|\tilde{k}|^2$ is therefore defined only away from that null set. The continuum small-spacing expansion for a fixed physical mode does not remove these high-frequency finite-grid modes.

Exact finite linearized constraint propagation uses commuting skew-adjoint derivatives and their Bianchi identity. It does not establish nonlinear pointwise closure or uniform agreement of the entire finite spectrum with continuum waves.

6 A small Leibniz defect is still a closure defect



Left: exact values for a five-site nonlinear bracket counterexample. Right: the centered Leibniz defect for $f = \sin x$ and $g = \cos 2x$ decreases with second-order behavior on smooth resolved data. These statements have different quantifiers.

For the finite pointwise derivative D, define

$$\mathfrak{L}_D(f, g) = D(fg) - (Df)g - f(Dg).$$

The scalar diffeomorphism model $\mathcal{D}_h[\xi] = h \sum_a \xi_a p_a (Dq)_a$ has bracket

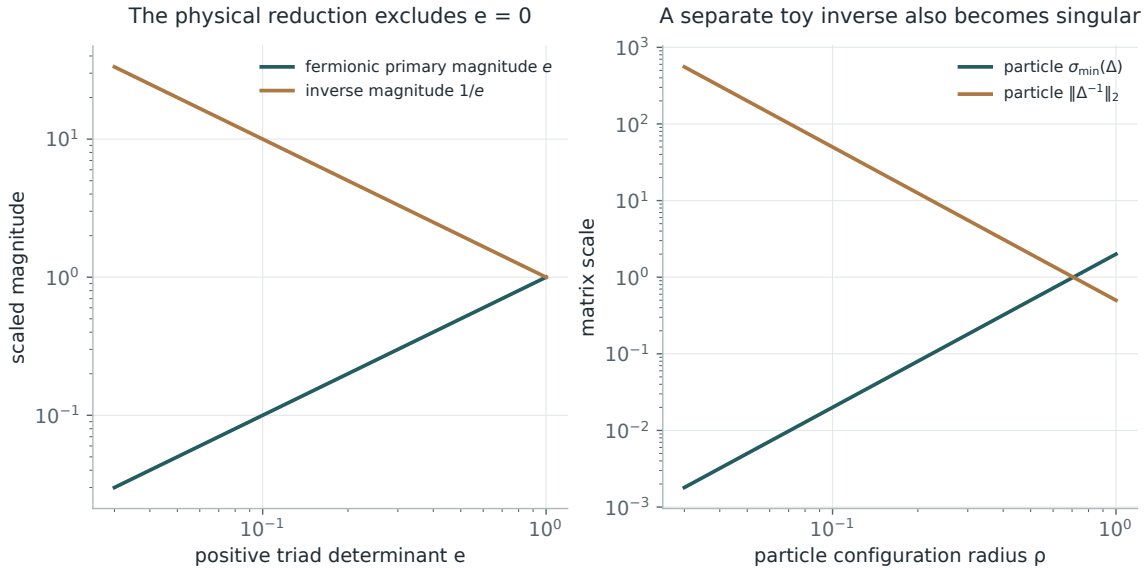
$$\{\mathcal{D}_h[\xi], \mathcal{D}_h[\eta]\} = \mathcal{D}_h[\xi D\eta - \eta D\xi] + h \sum_a p_a [\xi_a \mathfrak{L}_D(\eta, Dq)_a - \eta_a \mathfrak{L}_D(\xi, Dq)_a].$$

On five periodic sites with $h = 1$, take $q = e_0$, $p = \xi = e_1$ and $\eta = e_2$. Exact rational differentiation gives an actual bracket of zero, while the proposed closed right side is $-1/4$. Their difference is $1/4$. This is an explicit counterexample to that undeformed finite pointwise closure prescription.

For the smooth functions in the right panel, the centered defect is $O(h^2)$. Refinement therefore controls the particular error without making it identically zero on the full finite algebra. Conversely, a matrix inner derivation $\delta_K X = i[K, X]$ obeys the product rule exactly. The script checks one fully specified rational-matrix example.

Moving from the pointwise algebra to a matrix algebra removes this Leibniz obstruction, but does not supply every gravitational constraint, a continuum identification, or a globally invariant truncated band. Those additional objects and domain conditions belong to the monograph's resolved and transferred constructions.

7 Regular reduction has an inverse-matrix boundary



Two independent examples of inverse growth. Left: the magnitude of a fermionic primary block proportional to the triad determinant e and its inverse magnitude. Right: the Dirac matrix for a constrained-particle family with radius ρ. The two horizontal parameters are different objects.

The physical time-gauge reduction is restricted to the oriented nondegenerate sector $e = \det(e_a{}^i) > 0$. In the stated graded canonical convention,

$$\{\chi_A(x), \bar{\chi}_B(y)\}_{\mathrm{P}} = -ie(x)\delta_{AB}\delta^3(x-y).$$

The algebraic primary block is invertible when $e \neq 0$. The left panel displays its scalar magnitude and inverse magnitude after removing fixed identity factors. The plot approaches the excluded boundary; it supplies no extension through a degenerate triad.

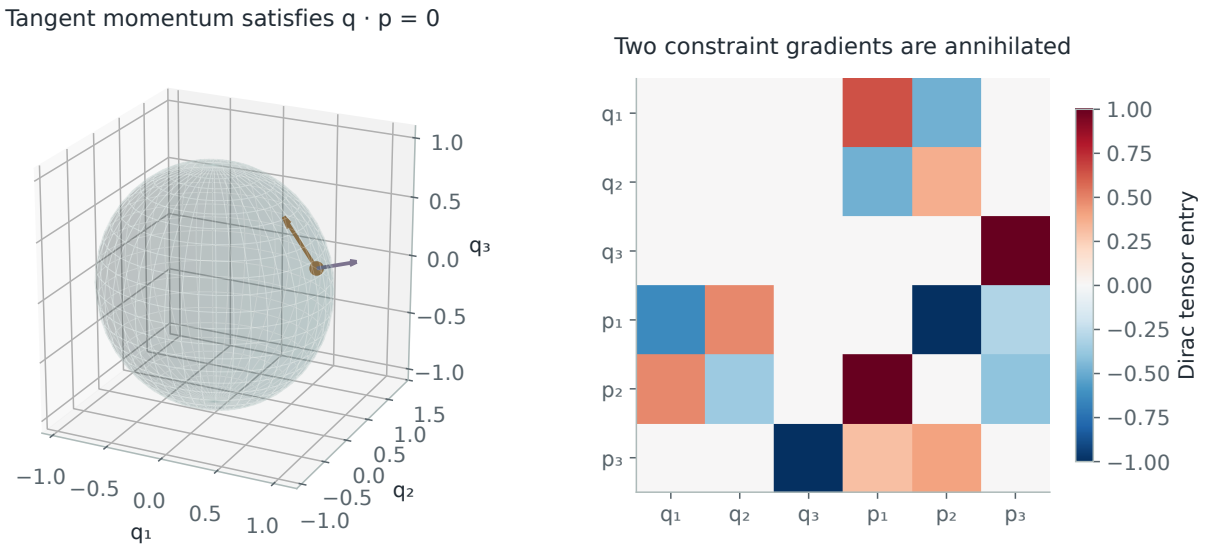
The right panel uses a separate mechanical example with constraints $q\cdot q - 1$ and $q\cdot p$. At $q = (\rho, 0, 0)$, $p = (0, 1, 0)$, its canonical Dirac matrix is

$$\Delta = GJG^T = \begin{pmatrix} 0 & 2\rho^2 \\ -2\rho^2 & 0 \end{pmatrix}, \qquad \sigma_{\min}(\Delta) = 2\rho^2.$$

Thus $\left\|\Delta^{-1}\right\|_2 = 1/(2\rho^2)$. This off-surface family diagnoses rank loss at the origin. Only $\rho = 1$ satisfies the unit-sphere constraint.

Both examples explain why local existence and stability statements need a lower bound for the relevant singular value. Neither parameter is the spherical horizon function. A zero outgoing expansion can occur in a regular Painlevé–Gullstrand chart, whereas a vanishing determinant or a singular gauge-fixing matrix is a boundary of the chosen reduction.

8 The Dirac tensor removes constraint directions



Left: a unit-sphere configuration with a tangent momentum arrow and a radial normal arrow. Right: the six-dimensional Dirac tensor at the rational point used in the exact check. Its two constraint-gradient rows are annihilated.

At a regular point, put $\Phi = (q \cdot q - 1, q \cdot p)^T$ and $G = \nabla\Phi$. The continuous Dirac tensor is

$$J_{\mathrm{D}} = J - JG^T(GJG^T)^{-1}GJ.$$

Direct algebra gives

$$J_{\mathrm{D}}^T = -J_{\mathrm{D}}, \qquad GJ_{\mathrm{D}} = 0.$$

Consequently, any vector field $\dot{z} = J_{\mathrm{D}}\nabla H$ is tangent to both constraint levels. This is a symplectic constraint construction, not an ordinary Euclidean orthogonal projection of an arbitrary spatial vector.

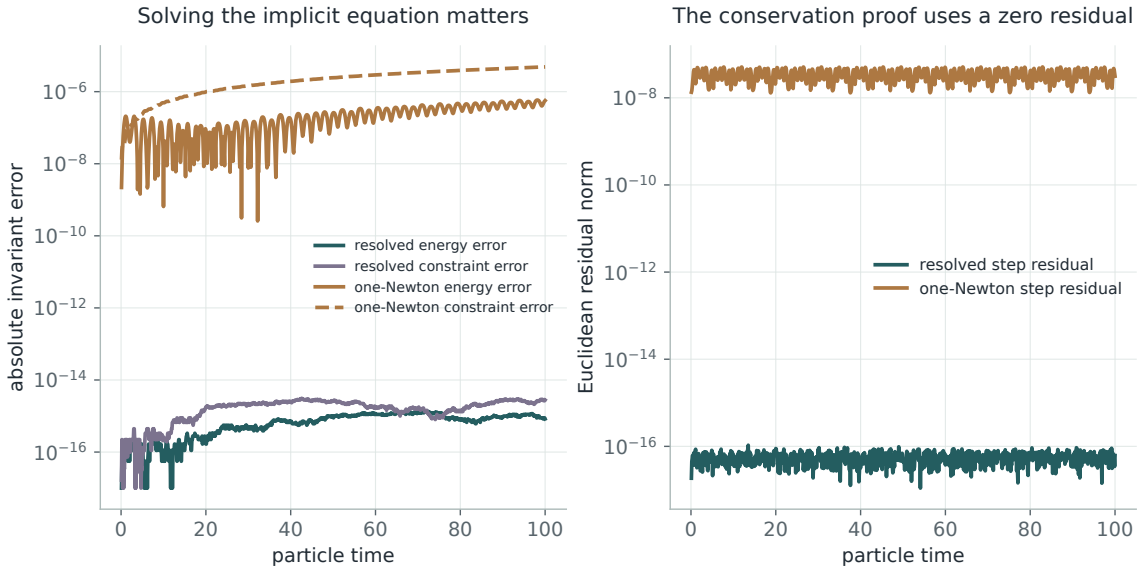
The exact benchmark takes

$$q = (3/5, 4/5, 0), \qquad p = (-4/5, 3/5, 1/2).$$

Here $q \cdot q = 1$ and $q \cdot p = 0$. Rational matrix arithmetic verifies skewness, annihilation, and rank four. The rank reduction from six to four removes the two second-class constraint directions. The heatmap shows the actual entries in the stated (q, p) basis.

The three-dimensional drawing displays the configuration sphere and a tangent momentum. It does not display the full six-dimensional phase space or the complete gravitational reduction. A finite gravity gauge patch uses its own constraints and gauge conditions, with a regular invertible Dirac matrix. For the endpoint discrete-gradient construction, ordinary derivatives in G are replaced by two-point discrete gradients.

9 The implicit equation must be solved



A fresh 1000-step quartic constrained-particle computation with $\Delta t = 0.1$. Fully resolved implicit steps are compared with a deliberately incomplete single Newton correction. The curves show actual floating errors and residuals, with values below 10^{-17} placed at the plotting floor.

A discrete gradient satisfies the exact chain rule

$$F(z_1) - F(z_0) = \overline{\nabla} F(z_0, z_1)^T (z_1 - z_0).$$

For the endpoint matrix G_Δ of constraint discrete gradients, define

$$J_{\mathrm{D}}^{\Delta} = J - J G_\Delta^T (G_\Delta J G_\Delta^T)^{-1} G_\Delta J.$$

An exactly solved implicit step obeys

$$z_1 - z_0 = \Delta t \, J_{\mathrm{D}}^{\Delta} \overline{\nabla} H.$$

Skewness and $G_\Delta J_{\mathrm{D}}^{\Delta} = 0$ then give exact preservation of H and Φ.

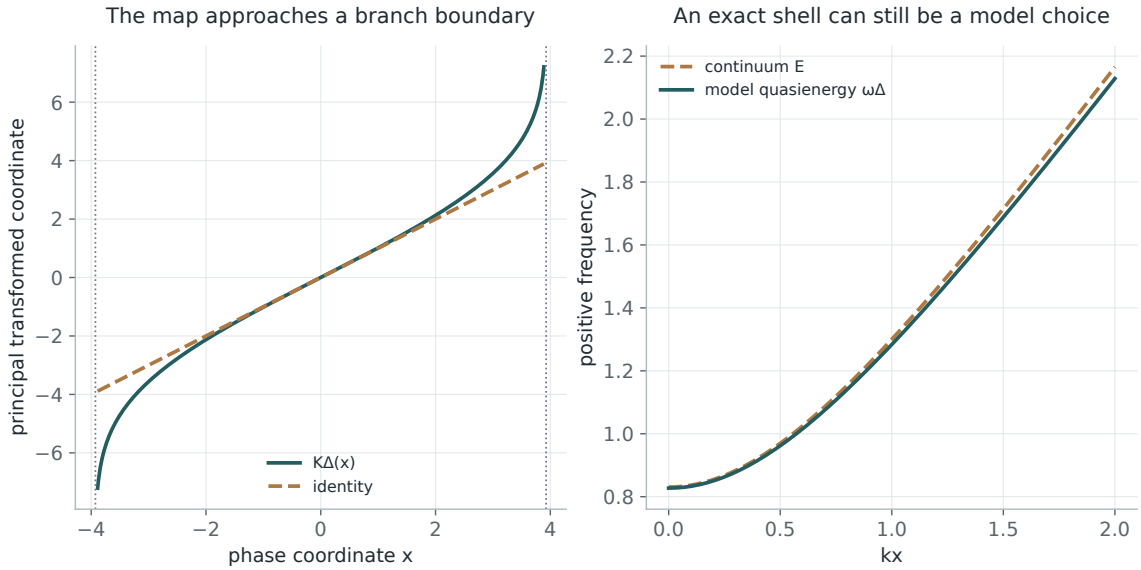
The test uses $H = |p|^2/2 + 0.2 \sum_i q_i^4/4$, $\Phi = (|q|^2 - 1, q \cdot p)$ and initial state $(q, p) = ((1, 0, 0), (0, 0.7, 0.2))$. Its polynomial average-vector-field gradients are evaluated explicitly. The resolved run uses SciPy's hybrid root solver; every step must satisfy the declared residual and regularity checks. The negative control starts from the same state but performs only one Newton correction per step.

If the retained solve residual is r, the invariant changes are

$$\Phi(z_1) - \Phi(z_0) = G_\Delta r, \qquad H(z_1) - H(z_0) = \overline{\nabla} H^T r.$$

Small numerical errors therefore support a bounded implementation check. They do not turn a floating solve into exact arithmetic. The particle experiment is a surrogate for the step algebra and is not a new Einstein–matter simulation.

10 An exact quasienergy formula retains a model choice



Left: the odd principal-coordinate map at $\Delta = 0.4$, approaching the cosine-zero boundaries. Right: the product-cosine model's positive quasienergy, compared with its continuum frequency for $m = 0.7$, $k_y = 0.4$, $k_z = 0.2$ and $0 \le k_x \le 2$.

The source's free product-cosine shell is

$$\cos(\Delta\omega_\Delta) = \cos(\Delta m) \prod_{i=1}^{3} \cos(\Delta k_i).$$

On the interval $|\Delta x| < \pi/2$, define

$$K_\Delta(x) = \operatorname{sgn}(x) \frac{\sqrt{-2 \log \cos(\Delta x)}}{\Delta}.$$

This odd increasing coordinate has inverse

$$K_\Delta^{-1}(P) = \frac{\operatorname{sgn}(P)}{\Delta} \arccos\left(e^{-\Delta^2 P^2/2}\right).$$

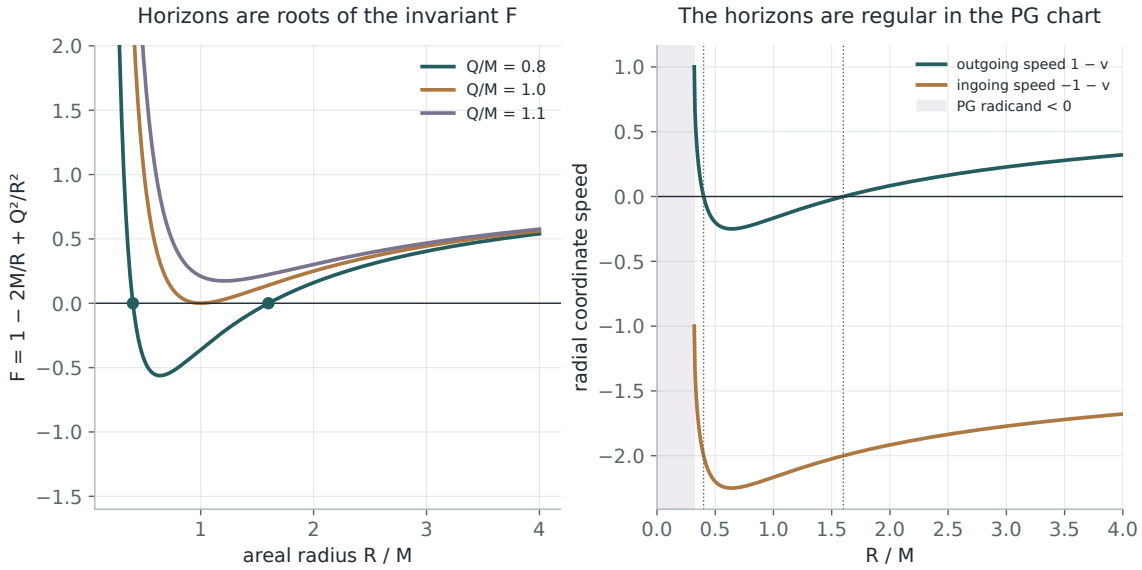
Within the principal positive-cosine domain, the product shell becomes

$$K_\Delta(\omega_\Delta)^2 - \sum_i K_\Delta(k_i)^2 = K_\Delta(m)^2.$$

The script checks the inverse on 501 inputs inside the branch and the shell on 401 momenta. Cosine zeros and other branches are excluded. A formula using $|\cos|$ also needs a sign label.

This exact relation belongs to the declared free model. It is not inferred from the ADM system. Factorization, branch consistency, continuum expansion, and empirical identification remain separate questions.

11 An electrovac horizon is regular in the PG chart



Left: the Reissner–Nordström horizon function for subextremal, extremal and superextremal charges, with $M = 1$. Right: radial null speeds for $Q = 0.8$ in the ingoing PG chart. The shaded interval has a negative PG radicand and is outside that real flat-slice chart.

For the charged spherical family,

$$F(R) = 1 - \frac{2M}{R} + \frac{Q^2}{R^2}, \qquad v(R) = \sqrt{\frac{2M}{R} - \frac{Q^2}{R^2}}.$$

The ingoing horizon-penetrating form is

$$ds^2 = -dt_{\mathrm{P}}^2 + (dR + v\,dt_{\mathrm{P}})^2 + R^2 d\Omega^2.$$

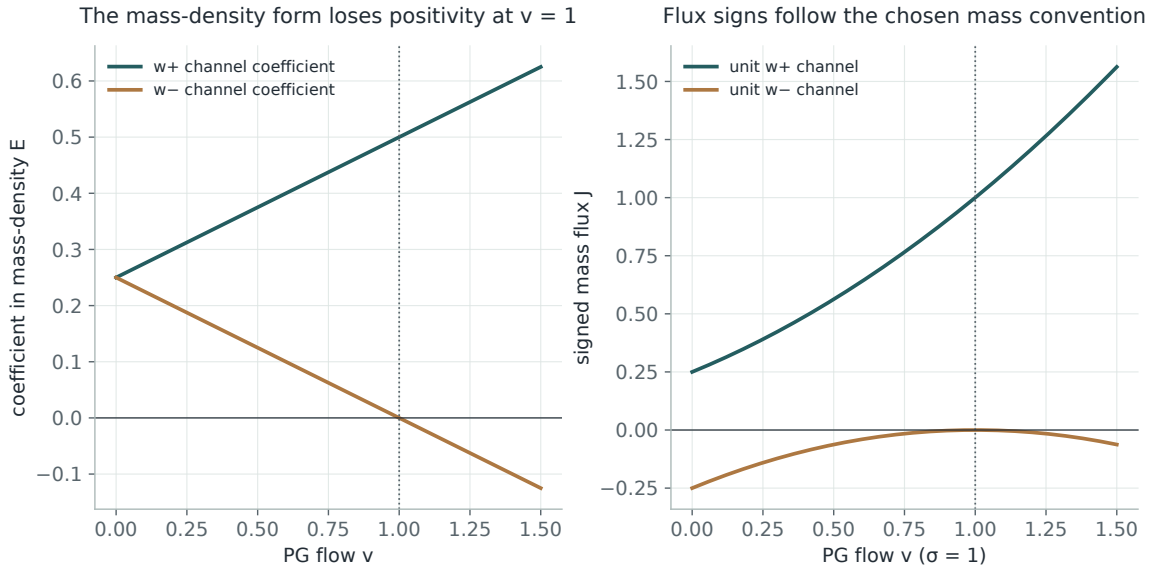
Its radial future null speeds are $dR/dt_{\mathrm{P}} = 1 - v$ and $-1 - v$. A horizon occurs at $F = 0$, equivalently $v = 1$ in this chart. Metric coefficients remain finite there.

For $|Q| < M$, the two radii are $R_{\pm} = M \pm \sqrt{M^2 - Q^2}$. At $|Q| = M$ they coincide. The plotted superextremal case has no real zero of the horizon polynomial. These statements concern the exact electrovac family, not matter-collapse outcomes.

The PG velocity is real only for $R \geq Q^2/(2M)$ when $M > 0$. Its turning boundary is separate from either horizon. Reconstructions that differentiate $\sqrt{2MR - Q^2}$ need a strict positive margin there. Likewise, the triangular mass-coordinate constraint transformation needs $R'/\Lambda \neq 0$ on its chosen monotone radial patch.

The atlas checks the horizon polynomial and plots this explicit family. It does not rerun a finite-range radial gravity slab or supply nonspherical stability. Conditioning near extremality and a uniform strong-field refinement estimate require the additional bounds recorded in the monograph.

12 Scalar mass density and mass flux have distinct signs



Coefficients of the two scalar characteristic channels as functions of the PG flow v. Left: the local mass-density quadratic form. Right: signed mass flux for one unit channel at a time, with $\sigma = 1$. These are pointwise analytic families, not an evolved scalar packet.

In the flat-slice scalar reduction,

$$m' = r^2\mathcal{E}, \qquad \dot{m} = r^2\mathcal{J},$$

where

$$\mathcal{E} = \tfrac{1}{2}(\Pi^2 + \Phi^2) + v\Pi\Phi, \qquad \mathcal{J} = \sigma(\Pi + v\Phi)(\Phi + v\Pi).$$

With $w_+ = \Pi + \Phi$ and $w_- = \Pi - \Phi$, exact expansion gives

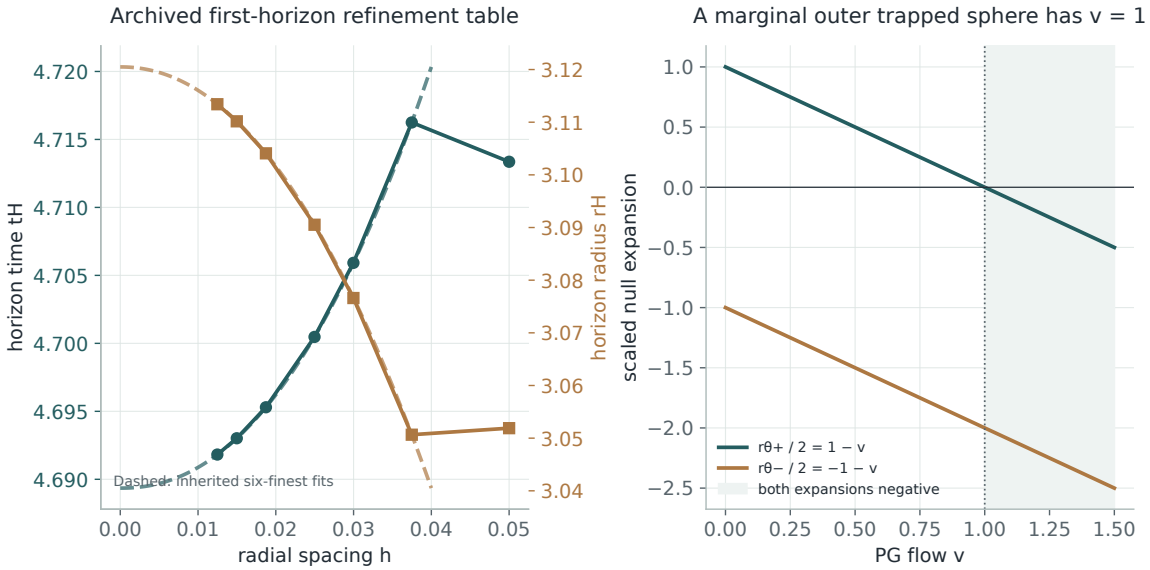
$$\mathcal{E} = \tfrac{1}{4}(1+v)w_+^2 + \tfrac{1}{4}(1-v)w_-^2,$$
$$\mathcal{J} = \tfrac{\sigma}{4}\left[(1+v)^2 w_+^2 - (1-v)^2 w_-^2\right].$$

For $0 \le v < 1$, both coefficients in the mass-density form are positive. At $v = 1$, the outgoing coefficient vanishes. For $v > 1$, this radial mass-density form is indefinite. A positive coercivity estimate valid outside a horizon cannot simply be carried unchanged into the trapped region.

The sign of $\mathcal{J}$ follows $\dot{m} = r^2\mathcal{J}$. Outside the horizon, the w_+ channel is ingoing and its positive flux increases enclosed mass, while w_- is outgoing. Inside, both coordinate characteristics are ingoing; the algebraic channel labels are retained. This mass-density form is not the local matter-energy norm. The compatibility identity $\partial_t(r^2\mathcal{E}) = \partial_r(r^2\mathcal{J})$ uses the scalar evolution and geometry equations together.

The two channel identities are checked symbolically for unrestricted symbolic Π, Φ, v, σ. Their plotted coefficients cover only $0 \le v \le 1.5$. No scalar profile or dynamical geometry is generated by this plate.

13 Inherited first-horizon data retain their provenance



Left: seven archived radial resolutions of the first marginally outer trapped sphere. Dashed curves are the inherited fits over the six finest runs. Time and radius use separate labeled axes. Right: analytic signs of the scaled null expansions as functions of v, independent of the archived numerical trajectory.

For the scalar PG chart, future radial null expansions can be normalized as

$$\theta_+ = \frac{2}{r}(1-v), \qquad \theta_- = -\frac{2}{r}(1+v).$$

For $r > 0$ and $v \geq 0$, $\theta_- < 0$. A marginally outer trapped sphere satisfies

$$\theta_+ = 0 \quad \Longleftrightarrow \quad v = 1 \quad \Longleftrightarrow \quad \frac{2m}{r} = 1.$$

For $v > 1$, both expansions are negative. This is a local trapped-sphere criterion, not a global event-horizon identification.

The inherited strong pulse has amplitude 0.10, center 10, width 1, and numerical radial domain $[0.05, 30]$. The finest archived grid has 2400 nodes and gives

$$t_{\mathrm{H}} = 4.69182094580, \quad r_{\mathrm{H}} = 3.11341154797, \quad m_{\mathrm{H}} = r_{\mathrm{H}}/2.$$

The recorded event time is interpolated between surrounding time steps. It precedes the stored last-step time, 4.69206969694. A field snapshot at that last step should not be silently relabeled as a field at the interpolated event.

The source scalar evolution stencil has radius two, while its mass and lapse projections are triangular radial prefix solves. The complete update is therefore not a strictly finite-range map. The archived result stops at first trapped-sphere formation and uses vacuum excision instead of a regular center. It does not determine critical collapse, singularity formation, cosmic censorship, or nonspherical stability.

14 Reproduce the finite assertions

From the editable source root, run

```
python tools/make_gravity_figures.py
python tools/build_documents.py Physics_Gravity_Visual_Atlas
```

The first command creates twelve PDF figures and twelve PNG previews. The second compiles this atlas with three XeLaTeX passes. The figure command uses NumPy, SciPy, SymPy and Matplotlib. It does not call Lean or any Einstein evolution script.

The authoritative record is `evidence/gravity/figure_checks.json`. It records the actual software dependencies, input checksums before and after the run, output checksums, numerical tolerances, sampled ranges, runtime, and every tested assertion. The particle check includes the endpoint Dirac-matrix lower bound and the maximum implicit solve residual. The archived event input is `evidence/gravity/inputs/archived_scalar_event_data.json`.

Each failed figure run is retained as a separate manifest. The append-only `figure_run_history.jsonl` records the succession of attempts. The successful record does not erase a parser failure or convert an incomplete run into a passing one. Figures are serialized in memory and written with atomic replacement, limiting the chance that an interrupted save leaves a zero-length result.

Arithmetic and its limits. Rational algebra and polynomial simplification check the exact finite identities. Numerical plots use double precision and stated tolerances; they are not interval-certified. SciPy's special functions and root solver remain dependencies of those numerical checks. The source derivations provide the general statements and their hypotheses. A successful figure manifest supplies evidence for the implemented examples, rather than a substitute proof of the whole gravity program.

Source correspondence. The plates follow the local balance, Gowdy evolution, centered constraint, quasienergy, regular reduction, electrovac, and scalar models developed in the gravity treatment. The verification companion and source concordance identify the exact source passages and computational records.

15 What the plates leave to the gravity program

The finite and symmetry-reduced checks make several useful boundaries concrete. The centered symbol has a null set that a projector must respect. Smooth convergence of a pointwise bracket defect does not yield exact nonlinear closure. An algebraic inverse can become poorly conditioned near a reduction boundary. The conservation argument for a discrete-gradient step requires the implicit equation and chain-rule gradients used in that argument.

Plate family	Evidence here	Additional obligation
Finite scalar work	Exact rational local balance	Coupled gravitational energy and physical identification
Gowdy and linear ADM	Analytic mode and finite symbol checks	Uniform evolution estimates beyond the reductions
Nonlinear defect	Exact finite negative control	A resolved or transferred algebra with all domains stated
Dirac-gradient step	Exact tensor algebra and bounded particle solve	Gravity gauge regularity, coercivity, and uniform tower bounds
Quasienergy model	Principal-branch finite identity	Independent physical identification of the model
Electrovac horizons	Exact spherical analytic family	Nonspherical dynamics and broader strong-field stability
Scalar trapped sphere	Archived refinement data, internal checks	Independent collapse reproduction and global continuum results

The resolved Fourier construction also needs enough support for every product and nested operation that appears in a bracket calculation. A fixed low-frequency band is generally not invariant under nonlinear multiplication. Its alias-free inequalities are domain conditions, and inverse-triad expressions are not automatically finite Fourier polynomials.

The classical BV–BFV and second-class reductions have their own hypotheses: regular derived moment maps, a tangent unreduced vector field, a nondegenerate triad, and an invertible gauge-fixing matrix where the local step is used. Finite illustrations do not remove these hypotheses. Quantum measure and anomaly questions, uniform infinite-refinement bounds, and a complete coupled radiative $3 + 1$ gravity solver remain separate obligations in the larger program.

The companion should therefore be read alongside this atlas. It records which statements are source derivations, which checks are executable finite evidence, which numerical data are inherited, and which formal building blocks apply to the gravitational constructions.

Adic and Polyphase Quantum Refinement: Visual Atlas

A visual atlas of normalized refinement and spectral structure

1 Read the address, normalization, and operator together

The adic structure in this volume belongs to the register addresses and their refinement. Amplitudes remain complex and probabilities remain ordinary nonnegative Born weights. A finite change of address, a Haar decomposition, a compatible Hamiltonian, and a Schur elimination are distinct operations. The twelve plates make their relations visible while preserving those distinctions.

The first four plates concern normalized descendants, child characters, cylinder averages, and finite unit coordinates. The next three concern compatible Hamiltonians, shell probabilities, and a local product-graph walk. The final five concern finite radial response, complete band folding, an infinite geometric tail example, prime spectral maps, and coherent number-residue packets.

Construction of the figures. All twelve figures were generated with Python and Matplotlib from small arrays or explicit scalar families. No inherited numerical table is reused. Exact modular, rational, and polynomial checks accompany floating matrix calculations. One scalar comparison uses sixty decimal digits. The computation record reports forty checks, including twelve elementary descendant normalization cases and an inventory check. These are checks on related implementations, rather than forty independent physical results.

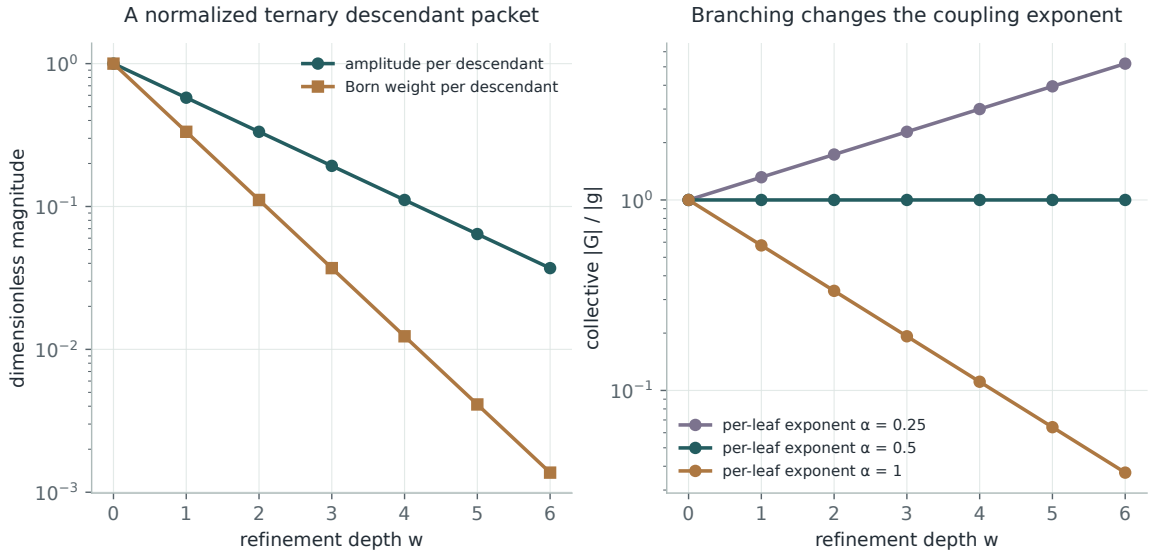
Conventions. At finite depth h, addresses are $x \in \mathbb{Z}/p^h\mathbb{Z}$ with counting-norm Hilbert space ℓ^2. The positive-exponent Fourier convention is

$$(F_p)_{qd} = p^{-1/2} e^{2\pi i q d/p}.$$

$E_r^{(h)}$ averages all addresses congruent modulo p^r. The detail projectors are $F_0^{(h)} = E_0^{(h)}$ and $F_r^{(h)} = E_r^{(h)} - E_{r-1}^{(h)}$ for $r \geq 1$. The symbol F_p with a prime subscript denotes the child Fourier matrix, while $F_r^{(h)}$ denotes a level projector. Tree couplings in these plates are real, and every resolvent evaluation lies in the upper half-plane.

How far the figures reach. The exact infinite geometric formulas in Plate 10 are a specialization of the manuscript's domain and truncation statements. Finite arrays do not prove the direct-limit theorem, self-adjointness, or every-depth convergence. The walk is a finite graph model, the radial chain is a chosen hidden-sector model, and the coherent packet plot is a standard oscillator calculation. The atlas reports no new Lean compilation or experimental calibration.

2 A descendant weight is not an interaction matrix element



Left: ternary descendant amplitudes and Born weights. Right: the collective visible-to-radial coupling for three independently chosen per-leaf exponents. All magnitudes are dimensionless and g is factored out.

After w normalized p-ary refinements, a parent basis vector becomes

$$|R_{p,w}\rangle = p^{-w/2} \sum_{\mathbf{d}\in\{0,\dots,p-1\}^w} |\mathbf{d}\rangle.$$

There are p^w orthonormal descendants, so the squared norm is one. Each descendant has amplitude magnitude $p^{-w/2}$ and probability p^{-w}. Using the probability as an amplitude would produce a norm defect. The script checks the probability sum exactly for $p = 2, 3$ and $w = 0, \dots, 5$.

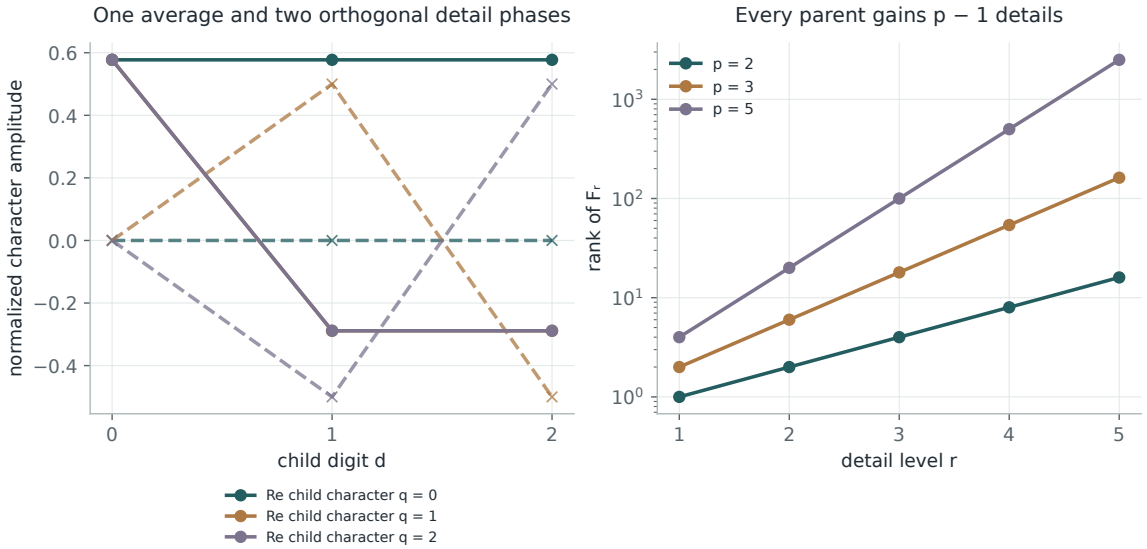
Choose a real per-leaf coupling $g\,p^{-\alpha w}$. Equal coupling to all descendants selects the radial packet and gives

$$G_{p,w} = \langle R_{p,w}|H|c\rangle = g\,p^{(1/2-\alpha)w}.$$

Thus $\alpha = 1/2$ preserves the collective strength, whereas $\alpha = 1$ makes the collective amplitude decay as $p^{-w/2}$. The rank-one self-energy then has numerator $|G_{p,w}|^2 = |g|^2 p^{-w}$.

Every descendant combination orthogonal to $R_{p,w}$ is dark to this equal-leaf source. This selection is a symmetry property of the chosen interaction. Haar normalization fixes the state geometry, but selecting α still specifies a Hamiltonian. A direct coupling to one normalized aggregate packet is another geometry and need not have the same exponent.

3 One child character is inherited and the others are details



Left: the real parts of the three normalized child characters are solid, and their imaginary parts are dashed. Right: the exact dimension of each detail level for $p = 2, 3, 5$. The finite matrix audit uses $h = 1, 2, 3$ with largest dimension 125.

The zero character is the normalized child average. The nonzero characters sum to zero and span the orthogonal detail complement. For a common parent, character orthogonality is

$$\frac{1}{p}\sum_{d=0}^{p-1} e^{2\pi i(q'-q)d/p} = \delta_{qq'}.$$

At depth h, the level-r wavelet for parent a and nonzero child character q is

$$\psi^{(h)}_{r,a,q}(x) = p^{-(h-r+1)/2} e^{2\pi i q d_r(x)/p} \mathbf{1}_{\{x \equiv a \bmod p^{r-1}\}}.$$

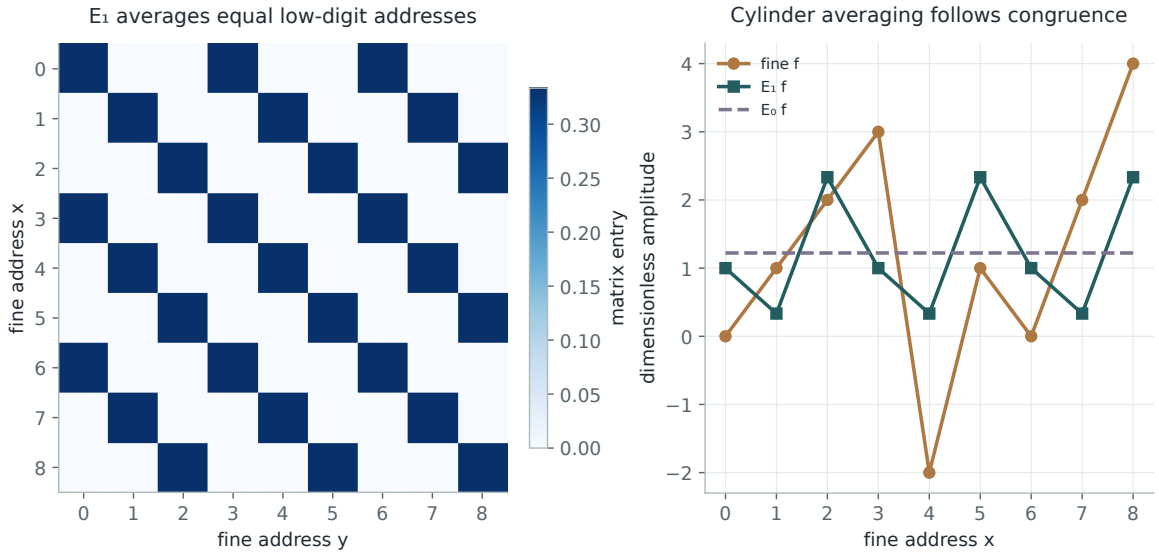
Its support has p^{h-r+1} points. The normalization factor makes its norm one. Disjoint parents are orthogonal, and the zero-mean condition gives orthogonality to older levels.

The independently assembled wavelet basis reconstructs

$$F^{(h)}_r = \sum_{a,q} |\psi^{(h)}_{r,a,q}\rangle\langle\psi^{(h)}_{r,a,q}|, \qquad \operatorname{rank} F^{(h)}_r = (p-1)p^{r-1}.$$

The script constructs cylinder differences separately and compares their entries with these wavelet outer products. It also checks the full basis Gram matrix and the trace ranks. The observed floating residuals concern nine finite cases. The uniform identity follows from the character-sum argument.

4 Cylinder averages follow the low digits



Left: the 9×9 ternary depth-two matrix $E_1^{(2)}$. Nonzero entries pair addresses with the same low ternary digit. Right: one fine vector and its depth-one and depth-zero averages.

The finite expectation is

$$(E_r^{(h)} f)(x) = p^{-(h-r)} \sum_{y \equiv x \pmod{p^r}} f(y).$$

Cylinder blocks are interleaved in the ordinary integer ordering. The matrix therefore need not look like contiguous blocks until the addresses are permuted. This is an ordering issue, rather than a failure of locality in the cylinder geometry.

Uniform refinement from p^h to p^{h+1} addresses is

$$(V_{p,h} f)(a + dp^h) = p^{-1/2} f(a).$$

It is an isometry in counting norm. At the next depth, its range is exactly the old-cylinder subspace:

$$V_{p,h}^* V_{p,h} = I, \qquad V_{p,h} V_{p,h}^* = E_h^{(h+1)}.$$

The cylinder expectations satisfy $E_r E_s = E_{\min(r,s)}$. The script checks every product of the three expectations at $p = 3, h = 2$ using exact rational matrices. It checks the refinement image separately in floating arithmetic.

In the infinite Haar probability space, a normalized finite basis vector corresponds to $p^{h/2} \mathbf{1}_{a+p^h \mathbb{Z}_p}$. That measure factor compensates the counting-norm replication factor. An unnormalized copy of amplitudes would describe a different map and would not preserve the finite state norm.

5 Unit coordinates retain the dyadic sign



Left: the eight units modulo 16 arranged as $(-1)^s 5^a$, with sign $s = 0, 1$ and phase $a = 0, 1, 2, 3$. Right: dimensions of the full register, phase image, and lost linear directions when sign is forgotten.

For $h \geq 0$, dyadic units have coordinates

$$(\mathbb{Z}/2^{h+2}\mathbb{Z})^\times \cong \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2^h\mathbb{Z}, \qquad (s, a) \longmapsto (-1)^s 5^a.$$

Odd-prime units have a separate torsion coordinate of size $p - 1$ and a principal coordinate of size p^h. These finite maps readdress a basis. When the labels are all retained, the corresponding Hilbert-space map is a permutation unitary.

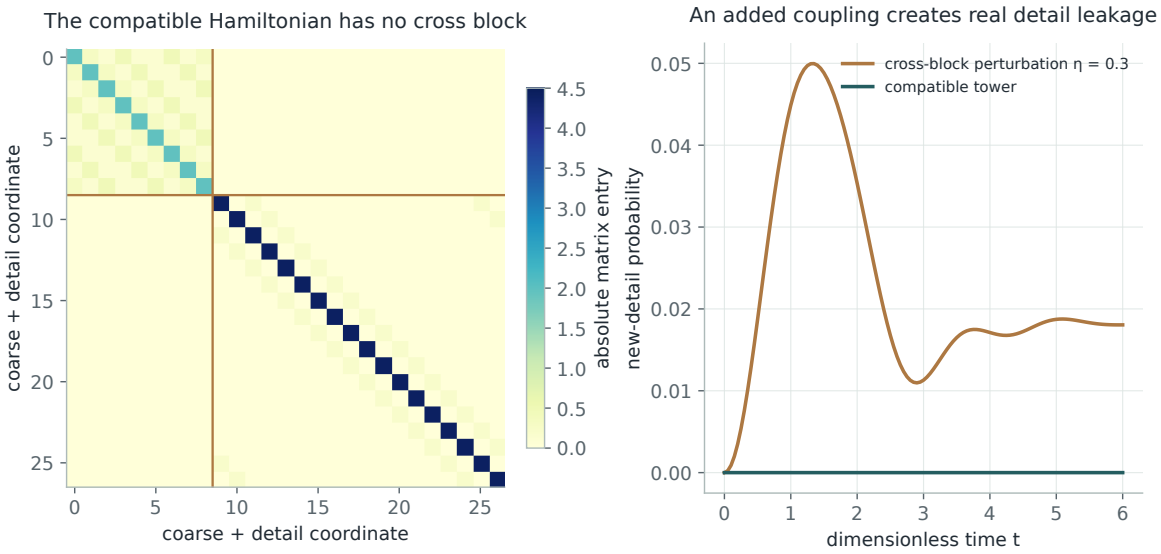
Deleting a coordinate has a different effect. The linear map

$$C_{\text{phase}}|\pm, a\rangle = |a\rangle$$

annihilates $|+, a\rangle - |-, a\rangle$. For the eight-dimensional example, its rank and kernel dimension are both four. The rank calculation is a negative control against interpreting coordinate deletion as a reversible change of basis. This linear map is not a partial trace or a normalized quantum measurement channel.

The fresh modular checks enumerate every unit and every coordinate-pair product for moduli 16, 9, and 25. These are exact integer checks on three chosen finite unit groups. The general decomposition uses the stated generators and the group-theoretic argument in the monograph. Sign-sensitive observables remain available because the dyadic sign factor is retained during principal-phase refinement.

6 Compatible dynamics have a vanishing cross block



Left: the absolute Hamiltonian entries in coarse-plus-new-detail coordinates for a ternary $9 \to 27$ tower. Right: new-detail occupation for a compatible Hamiltonian and a self-adjoint cross-block perturbation of size $\eta = 0.3$.

Let V be an isometry, $P = VV^*$, and $Q = I - P$. A self-adjoint pair is compatible when

$$H_f V = V H_c.$$

This is equivalent to $QH_fP = 0$ and $PH_fP = VH_cV^*$. Self-adjointness removes the opposite cross block as well. The old subspace is then reducing, so bounded functional calculus gives

$$e^{-itH_f}V = Ve^{-itH_c}.$$

At a nonreal energy, its cross-block Feshbach term also vanishes.

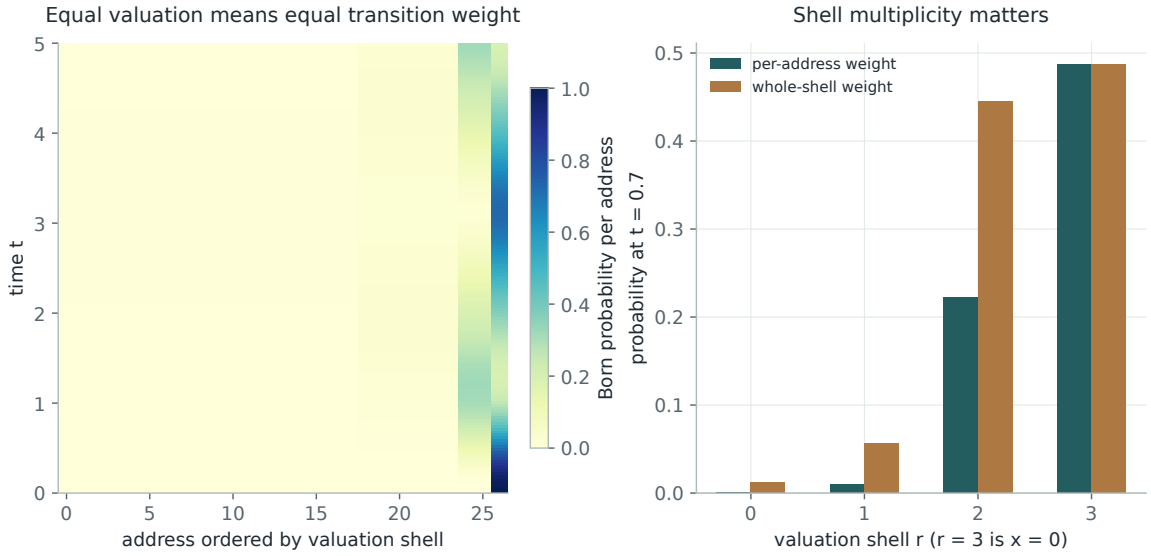
The displayed family uses

$$H_h = \sum_{r=0}^{h} \epsilon_r F_r^{(h)} + 0.2(2I - S_h - S_h^*), \qquad \epsilon_0 = 0, \quad \epsilon_r = 2^{r-1}\ (r \geq 1).$$

The script constructs the cyclic shifts and cylinder projectors independently at the two depths. It checks Hamiltonian, unitary, leakage, and self-energy relations within stated tolerances.

For the negative control, $u = Ve_0$ and v is a normalized new-detail wavelet. Adding $\eta(|u\rangle\langle v| + |v\rangle\langle u|)$ produces cross-block norm η and visible detail occupation. Self-adjointness alone therefore does not imply refinement compatibility. The zero leakage in the compatible curve is a structural condition, not a universal property of arbitrary refinements.

7 Hierarchical evolution is uniform within valuation shells



Left: per-address probabilities for the pure ternary hierarchy, ordered by valuation shell, over $0 \le t \le 5$. Right: per-address and whole-shell probabilities at $t = 0.7$. The final shell contains only the source address.

For orthogonal level projectors and real energies,

$$H_h = \sum_{r=0}^{h} \epsilon_r F_r^{(h)}, \qquad U_h(t) = \sum_{r=0}^{h} e^{-it\epsilon_r} F_r^{(h)}.$$

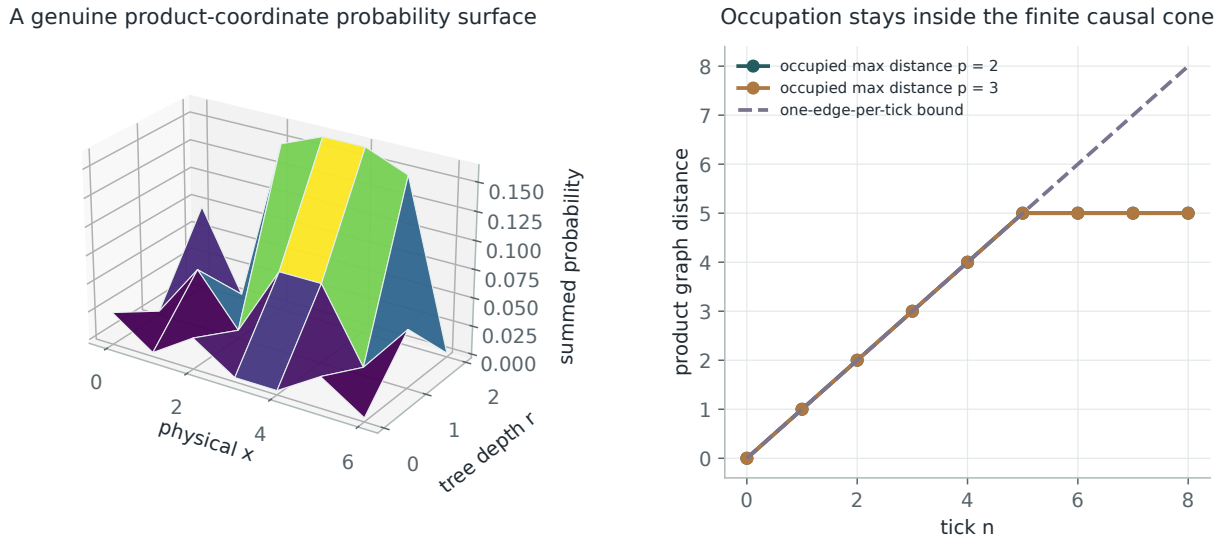
Every expectation kernel is a normalized congruence indicator. The resulting propagator kernel therefore depends on the valuation of the address difference, truncated at the register depth. Equal valuations have equal amplitudes in this pure hierarchical model.

The plate uses $p = 3$, $h = 3$, and energies $(0, 1, 2, 4)$, starting at address zero. Each time slice is constructed directly from the projector sum. Probability normalization and the spread within every valuation shell are checked at all 151 plotted times.

Uniformity is per address. The probability of an entire shell also includes its multiplicity. For $r < h$, the shell $v_p(x) = r$ contains $(p-1)p^{h-r-1}$ addresses; the diagonal source is a singleton. The two right-panel bar sets therefore describe different observables.

Adding a cyclic Laplacian preserves the compatible tower used in the previous plate, but its kernel need not retain this pure hierarchical valuation-shell form. Cross-prime couplings can likewise change shell symmetries. The current plot isolates the spectral-projector mechanism and does not infer shell isotropy for all compatible Hamiltonians.

8 A product-graph walk propagates one edge per tick



Left: a genuine three-dimensional probability surface at tick six on $C_7 \square T_{2,2}$, summed over all tree vertices of each depth. The connecting surface guides the eye between discrete samples. Right: occupied product distances in two finite walks and the one-edge-per-tick bound.

The vertices are a physical cycle address and a rooted-tree prefix. A product edge changes exactly one of them. At each vertex, the Grover coin acts on outgoing arcs, and the flip-flop shift reverses an arc:

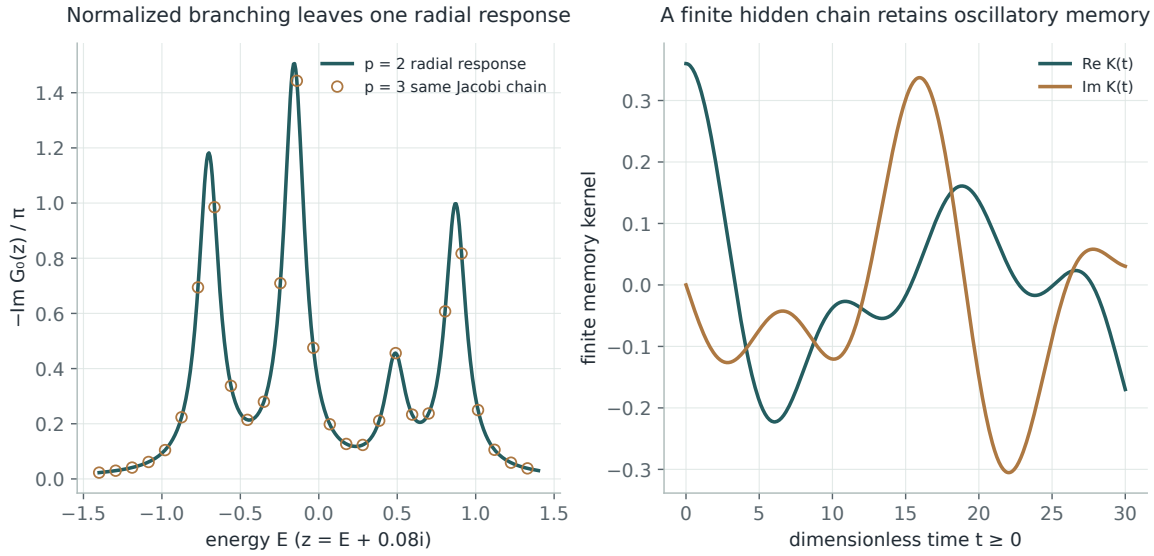
$$U = SC, \qquad C_v = 2|+_{\deg v}\rangle\langle +_{\deg v}| - I.$$

The coin is local and unitary, while $S^2 = I$. Starting at one vertex, occupation remains inside the product ball $d_C + d_T \le n$ after n ticks.

The fresh calculation runs eight ticks on $C_7 \square T_{p,2}$ for $p = 2, 3$. It checks probability, the local incoming-minus-outgoing continuity identity, and the probability outside the exact cone. The latter is zero in the constructed finite arrays. The left surface is an aggregate of graph coordinates and probabilities, not a metric embedding of physical spacetime.

Finite-depth locality does not supply a depth-refinement identity. Under canonical inclusion of old arcs, extending the tree changes an old leaf's degree from three to five in the binary example. Its Grover diagonal entry changes from $-1/3$ to $-3/5$. This explicit negative control prevents transferring the compatible Haar Hamiltonian theorem to this walk without a separately constructed intertwiner.

9 Branch normalization isolates a prime-independent radial chain



Left: the root resolvent response of one finite radial chain, with $z = E + 0.08i$. The ternary markers use the same radial matrix as the binary curve. Right: the real and imaginary parts of the exact finite hidden-chain memory kernel.

At tree level r, give each edge coupling $\gamma_r/\sqrt{p}$ and each vertex onsite energy ϵ_r. The normalized level state is

$$|r\rangle = p^{-r/2} \sum_{|v|=r} |v\rangle.$$

Its neighboring radial matrix element is γ_r, since the branching enhancement cancels the edge normalization. The invariant radial restriction is a Jacobi chain J with diagonal ϵ_r and off-diagonal γ_r.

For a nonreal z, repeated Schur elimination gives

$$\Sigma_h = 0, \qquad \Sigma_r(z) = \frac{\gamma_r^2}{z - \epsilon_{r+1} - \Sigma_{r+1}(z)}, \quad G_0(z) = \frac{1}{z - \epsilon_0 - \Sigma_0(z)}.$$

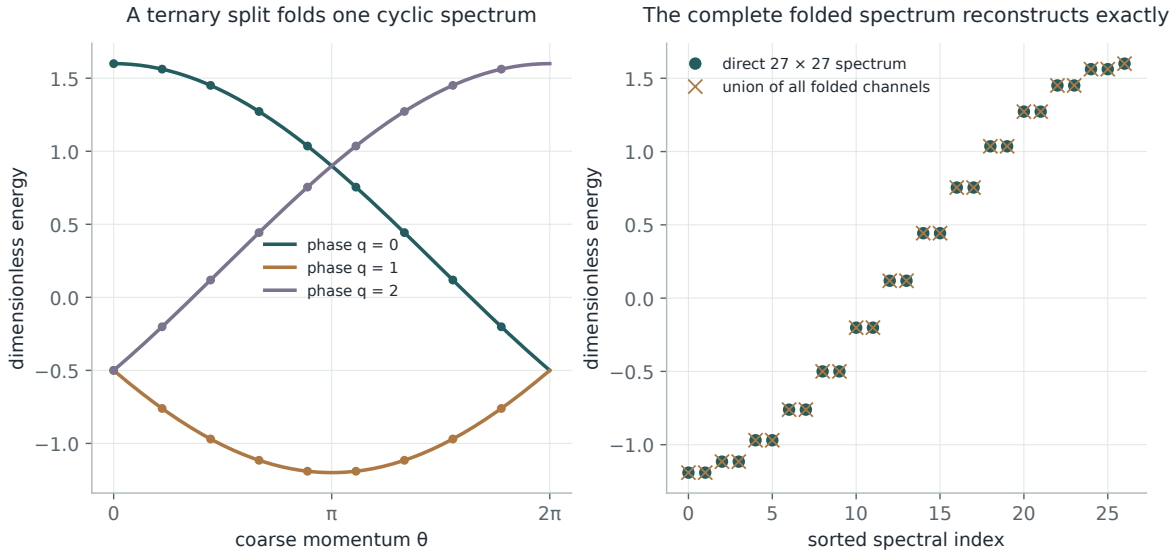
The script compares the full binary and ternary tree root inverses with the Jacobi inverse and the continued fraction. The depth-three examples have 15 and 40 tree vertices. Prime dependence survives in dark-sector multiplicities.

For a visible source coupled to the first tail site, the initially empty hidden sector contributes the causal convolution kernel

$$K(t) = \gamma_0^2 \langle 0 | e^{-itJ_{\text{tail}}} | 0 \rangle, \qquad t \geq 0.$$

A finite tail has persistent spectral oscillations. This plot does not establish irreversible damping, an infinite-bath limit, or an experimentally selected hidden sector. The same finite algebra can support those questions once the additional model and limit are specified.

10 Band folding preserves the spectrum when every channel is kept



Left: the three folded bands for $\mu = 0.2$, $\kappa = 0.7$, with markers at the nine coarse momenta of a 27-site cycle. Right: the sorted direct matrix spectrum and the complete folded-band union.

Standard polyphase analysis sends $|pk + r\rangle$ to $|r\rangle \otimes |k\rangle$. Adic analysis splits the newest digit instead. Digit reversal converts one ordering to the other. Both are permutation unitaries while all phase channels are retained.

For the circulant Hamiltonian

$$H = \mu I + \kappa(S + S^*),$$

the folded bands at coarse momentum θ are

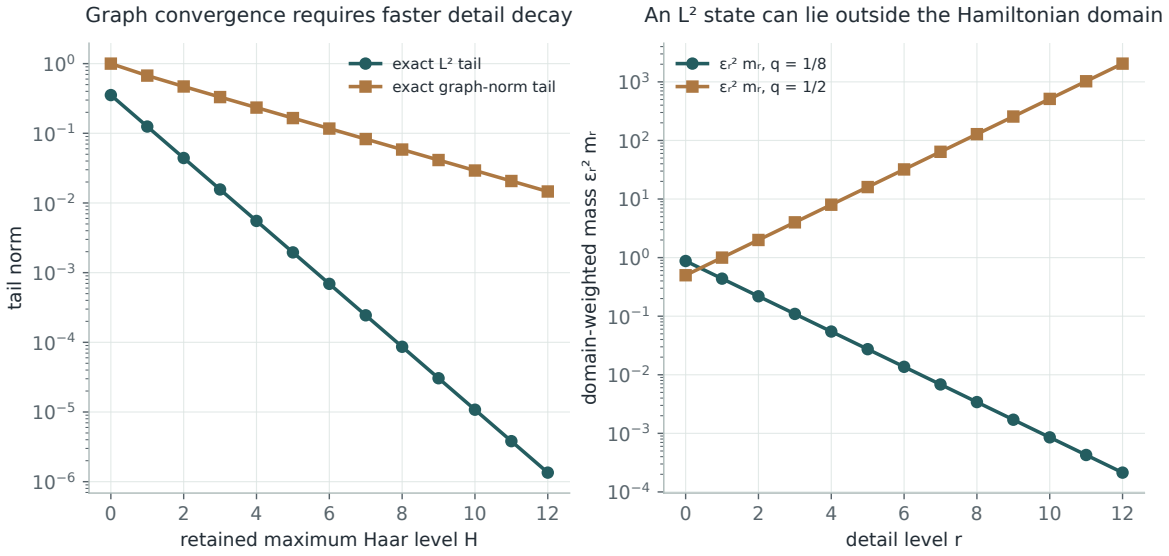
$$E_q(\theta) = \mu + 2\kappa \cos\left(\frac{\theta + 2\pi q}{p}\right), \qquad q = 0, \ldots, p-1.$$

At a finite register depth, the original fine index has the unique decomposition $k = j + qm$. Consequently the union of all folded bands reproduces the fine cyclic spectrum, including multiplicity.

The fresh ternary check compares these scalar samples with the eigenvalues of the independently assembled 27×27 cyclic matrix. Keeping only one folded channel would leave nine spectral entries and discard eighteen. This dimension count is a negative control against presenting a channel truncation as an invertible coordinate change.

Schur elimination can retain the low-channel resolvent through an energy-dependent self-energy when the hidden inverse exists. That reduced operator is a different operation from keeping the full polyphase symbol. Near a hidden pole, the eliminated representation requires the stated resolvent restrictions and may become singular.

11 An infinite tail can converge in norm but fail in graph norm



Left: exact infinite L^2 and graph-norm tail formulas for $q = 1/8$ and $\epsilon_r = 2^r$. Right: domain-weighted level masses for this example and for $q = 1/2$. The curves do not hide a finite depth ceiling.

Choose one unit vector in each infinite Haar level and assign squared coefficient

$$m_r = (1-q)q^r, \qquad r \geq 0, \quad 0 < q < 1.$$

These masses sum to one. With hierarchical energies $\epsilon_r = 2^r$, membership in the Hamiltonian domain requires

$$\sum_{r\geq 0} \epsilon_r^2 m_r < \infty, \qquad \text{equivalently } 4q < 1.$$

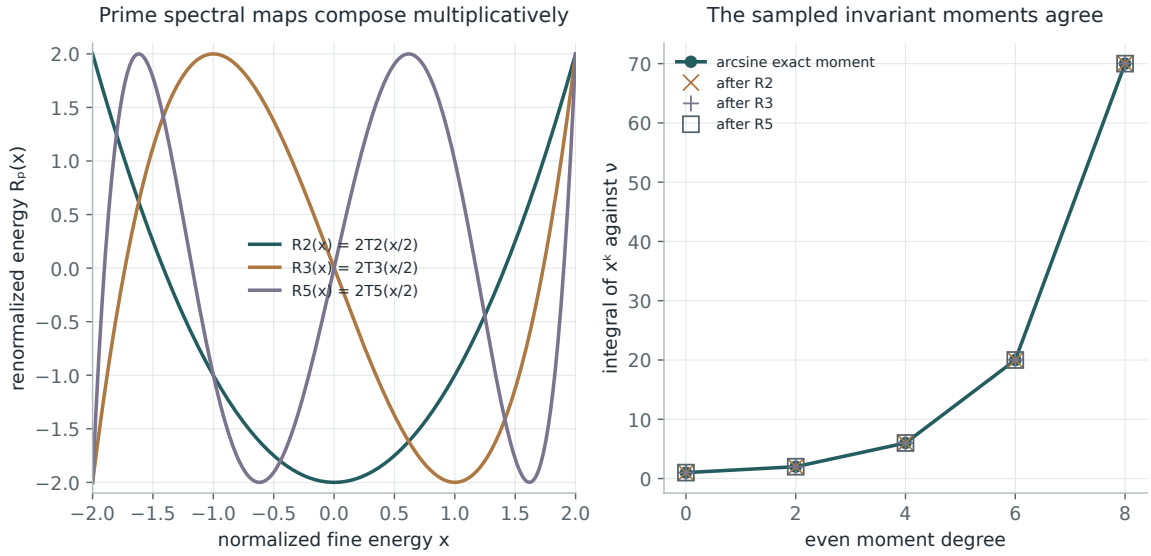
For $q = 1/8$, the exact tails after retaining level H are

$$\|(I - E_H)f\|_2^2 = q^{H+1},$$
$$\|(I - E_H)f\|_{\mathcal{D}(H)}^2 = q^{H+1} + \frac{(1-q)(4q)^{H+1}}{1-4q},$$

where the graph norm is $\|f\|_{\mathcal{D}(H)}^2 = \|f\|_2^2 + \|Hf\|_2^2$. The plot evaluates these closed infinite geometric sums. A separate 100-term check agrees with them within the recorded tolerance.

For $q = 1/2$, the state still has unit L^2 norm, but its domain-weighted summand is 2^{r-1} and fails to tend to zero. The norm convergence of cylinder expectations therefore cannot be upgraded to convergence of Hf for this vector. General domain statements need the weighted-square-summability hypothesis, even though unitary evolution is defined on every L^2 vector by spectral calculus.

12 Prime spectral maps preserve an arcsine law



Left: the normalized maps $R_p(x) = 2T_p(x/2)$ for $p = 2, 3, 5$. Right: exact even arcsine moments and the same moments after prime maps, evaluated with 4096 angle nodes.

On the free spectral interval, write $x = 2\cos\theta$. Then

$$R_p(2\cos\theta) = 2\cos(p\theta), \qquad R_p \circ R_q = R_{pq}.$$

The script checks the 2- and 3-step composition as an exact polynomial identity. For example, $R_2(x) = x^2 - 2$ and $R_3(x) = x^3 - 3x$.

The arcsine probability measure is

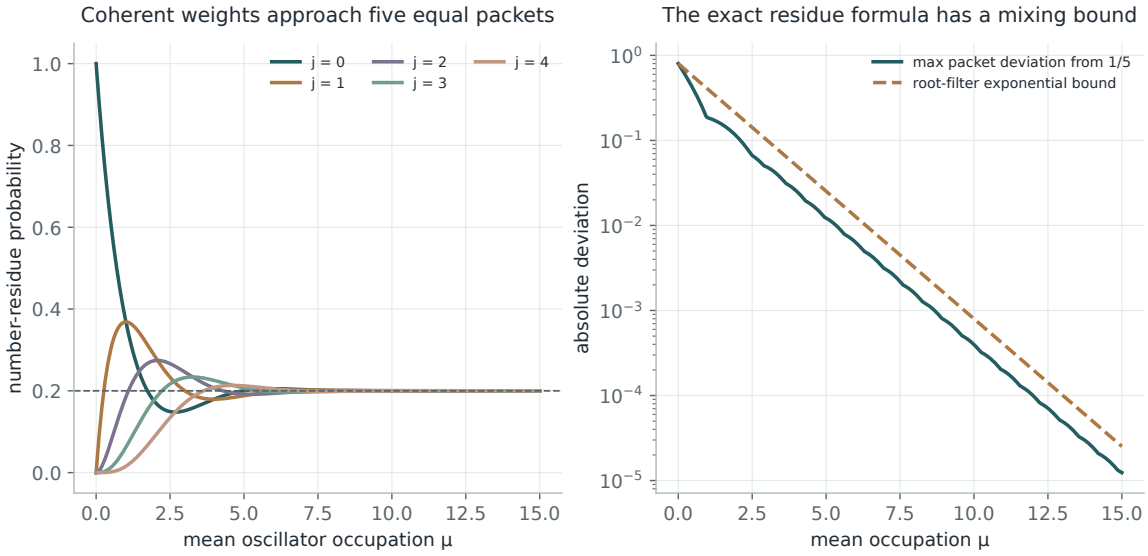
$$d\nu(x) = \frac{dx}{\pi\sqrt{4 - x^2}}, \qquad -2 < x < 2.$$

It is the cosine pushforward of uniform angle. Multiplying the angle by an integer preserves that pushforward distribution, which proves $(R_p)_*\nu = \nu$. The fresh numerical comparison tests moments of degrees zero through eight. Their exact values are zero for odd degree and $\binom{2n}{n}$ for degree $2n$.

In the Chebyshev orthonormal basis, the Koopman isometry sends mode n to mode pn. Its adjoint retains divisible modes and removes the others. This is an isometry with a proper range, rather than an invertible spectral map. The plotted maps have several inverse branches. Recovering an eliminated fine state therefore requires the missing channel information.

The normalized polynomial map extends across the spectral interval, while a particular Schur formula can have a hidden Dirichlet pole there. Polynomial continuation does not make the hidden resolvent invertible at that pole. The figure displays the normalized map, not an unrestricted inverse for every eliminated operator.

13 Hypergeometric lanes partition coherent number probabilities



Left: probabilities in the five number-residue packets of a coherent oscillator state. Right: the maximum deviation from $1/5$ and its root-filter exponential bound. The curves use 301 values of $0 \le \mu \le 15$.

For $N \ge 2$, the exponential lane is

$$E_{N,j}(z) = \sum_{q \ge 0} \frac{z^{Nq+j}}{(Nq+j)!} = \frac{1}{N} \sum_{\ell=0}^{N-1} \omega_N^{-j\ell} e^{\omega_N^{\ell} z}.$$

Splitting the factorial into N step lanes gives the same function in generalized hypergeometric form,

$$E_{N,j}(z) = \frac{z^j}{j!}\, {}_0F_{N-1}\left(;\mathcal{A}_{N,j};\frac{z^N}{N^N}\right),$$

where $\mathcal{A}_{N,j}$ is the multiset of $(j+r)/N$, $1 \le r \le N$, with the one parameter equal to 1 removed.

A coherent state has Poisson number probabilities with mean $\mu = |\alpha|^2$. Its residue weight is $p_{N,j}(\mu) = e^{-\mu} E_{N,j}(\mu)$, so positivity follows from the series and normalization from lane reconstruction. For fixed N,

$$\left| p_{N,j}(\mu) - \frac{1}{N} \right| \le \frac{N-1}{N} e^{-\mu[1-\cos(2\pi/N)]}.$$

The fresh calculation compares root sums, sparse factorial series, and hypergeometric expressions for $N = 3, 5$ at sixty decimal digits. It also checks the plotted normalization, positivity, and bound. These oscillator spectral packets are actual orthogonal number projectors. Individual functions $E_{N,j}(z)$ are analytic lanes, and should not be called Hilbert-space projectors or assumed bounded when evaluated on an arbitrary unbounded Hamiltonian.

14 Reproduce the figures and inspect the evidence

The editable generator is `tools/make_adic_figures.py`. From the source root, run

```
python tools/make_adic_figures.py
```

It writes twelve PDF plates and twelve PNG previews in `source/adic_figures/`. It writes the computation record to `evidence/adic/figure_checks.json` and appends a small run-history entry. If a calculation fails, the script retains the failed manifest and the exact script snapshot, reports the failure on stderr, and exits with a nonzero status. A partial plate set is not a successful atlas run.

What the record contains. The record pins the executed source, each declared output, the Python and numerical-library versions, the UTC start and elapsed runtime, and the arithmetic and finite bounds of every check. The current run used no random sampling. There was no network request or inherited numerical-data dependency. The source correspondence map separately pins the archived chapter and script members from which the mathematical topics were developed.

What the checks mean. Exact finite modular and rational operations decide the specified finite assertions. Symbolic polynomial checks certify the displayed polynomial expressions through the installed symbolic library. Matrix and scalar comparisons are conditional on their recorded tolerances. The high-precision scalar calculation is not interval certified. Library versions are recorded but the environment is not locked, and no operating-system memory or CPU cap was imposed.

The independent constructions matter: Haar wavelet outer products are compared with separately assembled cylinder differences, folded scalar bands with a cyclic matrix spectrum, full tree inverses with radial inverses and continued fractions, and root filters with factorial and hypergeometric forms. Negative controls expose the failure of probability-as-amplitude normalization, sign deletion, cross-block compatibility, old-arc Grover refinement, single-channel spectral reconstruction, and unrestricted graph-domain membership.

Formal provenance. Python checks do not replace a Lean kernel receipt. The companion maps the archived formal support and its limits. This atlas runs no Lean compilation and does not reinterpret historical build logs as new formal certification.

15 Keep the exact interface with its hypotheses

The visual argument passes through several interfaces. Each one carries its own condition.

Interface	Condition preserved by the calculation
Address factorization	Every coordinate is retained for a permutation unitary.
Haar refinement	Counting amplitudes include $p^{-1/2}$ per new digit.
Compatible dynamics	The old subspace reduces the self-adjoint fine Hamiltonian.
Graph walk locality	The coin is vertex local and the shift crosses one edge.
Radial compression	Level symmetry and branch-normalized couplings are specified.
Schur reduction	The relevant hidden block has an inverse at the chosen energy.
Infinite Hamiltonian	The vector lies in the weighted spectral domain when Hf is used.
Analytic residue calculus	Bounded operators admit the norm series; unbounded ones need domains or spectral cutoffs.

The direct-limit Hilbert identification and the self-adjoint spectral construction are analytical results developed in the monograph. The finite matrices illustrate their concrete coordinates. A forty-check success record therefore supports these displayed examples and their implementation, while the proofs retain responsibility for the uniform claims.

The remaining physical interpretation is also explicit. A tree edge is a refinement adjacency, a finite memory kernel belongs to the selected hidden chain, and a number-residue probability belongs to an oscillator observable. The calculations provide reproducible mathematical coordinates for those models. A microscopic action, calibrated interaction strengths, an interacting unbounded theory, or an observable that selects a particular digit ordering would require additional work.

> **What we have established**
>
> The atlas establishes the stated finite and explicit-family checks, with their input and output identities recorded. It preserves the difference between a unitary change of coordinates, a compatible tower, a conditional elimination, and an infinite operator domain. Those boundaries let the figures explain the mathematics without claiming a completed physical theory.

PART VI

Verification and computational evidence

A proof, a compiler certificate, a numerical check, and a physical interpretation answer different questions. The following appendices identify their exact scope. Source and toolchain identifiers appear here only where they are needed to inspect or reproduce the evidence.

IN THIS PART

Verification Scope and Source Conventions 474

Foundations Formalization Companion 477

Dirac Formalization Companion 491

Gauge Formalization Companion 508

Gravity Verification Companion 522

Adic Verification Companion 540

Fiber Model: Computational Evidence 557

Cosmology: Reproducibility and Data 564

Mathematics, evidence, and interpretation

The scope of a proof

A theorem asserts what follows from its stated hypotheses. A conditional theorem remains a theorem under those hypotheses. Numerical evidence concerns the declared inputs, tolerances, and implementation. A physical model adds an identification of mathematical objects with quantities in nature. These roles are kept separate throughout the volume.

The free Dirac construction

The formal record covers mode algebra, completed coefficient and torus evolution, the maximal Hamiltonian, its exact first-order Fourier domain, physical self-adjointness, and the full strong generator. The centered extension proves graph-domain convergence, actual nonreal resolvent inverses and strong limits, and quadratic estimates under a finite third-order Fourier moment. The abstract spectral-functional-calculus identity remains a written deduction without a separate formal API linkage.

The comparison spaces matter

The centered operators act on the same infinite torus Fourier index space. A changing finite spatial grid additionally needs sampling, interpolation, and aliasing control. A background coupling changes the operator-domain problem. A dynamical gauge or gravitational field requires further evolution and constraint arguments.

Reading the technical appendices

The detailed companions identify exact declarations, hypotheses, source records, computations, and missing corroboration. A recorded compiler pass is not a new independent compilation. A checksum validates identity relative to a reference file, not the physical assumptions or the truth of every surrounding interpretation.

VERIFICATION AND SOURCES

Formal and computational source identity

Pinned formal source

The preserved operator baseline is Komorebi Physics Lean v0.49.0, Checked Block 49. The unchanged Relativistic3DGeneratorBridge is extended by centered-convergence, modal-perturbation, centered-rates, and centered-resolvent modules in the KomorebiPhysics namespace. All use Lean 4.34.0-rc2 and Mathlib revision aeb72f1e9981b6d35b426b16dca7759be3e7ba06. Exact source identities belong to the reproduction record, not the mathematical chapter names.

Current and inherited compiler evidence

The current run freshly compiled twenty-six local modules and audited 209 environment declarations, including generated auxiliaries: 162 from the four new modules and 47 from the unchanged generator. Five module kernel replays and twenty-one positive or rejection controls passed. Source, object and relevant runtime identities remained stable. The sixty-four command records retain their individual outcomes. The compiler and external dependencies were pinned and reused, not rebuilt from source.

Known evidence limitations

The gauge records contain a cumulative-build log with a digest discrepancy and a Git exit-code-128 failure. Six referenced logs for two later source checkpoints are absent. Source identity checks do not turn those gaps into successful builds. The detailed gauge companion identifies the affected records. The generic aggregate Lean root also does not establish the full formal endpoint by itself. The named module and its exact import closure are the relevant target.

Numerical evidence

The preceding publication checkpoint reran five supplied suites: finite Fourier generator tests, exact Clifford and edge checks, finite adic identities, Lorentz-sector BRST parity, and the covariant backreaction benchmark. Their retained logs and input hashes preserve the energy-availability obstruction. The current formal extension does not rerun those computations or every cosmology, electrovac, and collapse simulation. External-data and historical reproduction limits remain explicit.

VERIFICATION AND SOURCES

Physical assumptions and source conventions

Energy availability controls the matter interpretation

For the specified rapid-crossing benchmark, the unprojected spectrum exceeds the available crossing energy by about 2.63×10^{20}. A common normalization factor does not solve the coupled production equations. The cosmology calculations therefore propagate a prescribed distribution. They are conditional model results, not a demonstrated microscopic production prediction.

Gravity and arithmetic refinement

The gravitational results retain their chart, nondegeneracy, boundary, spectral-support, and numerical hypotheses. The classical master equation does not define a quantum measure or settle anomalies. Arithmetic refinement supplies written finite and infinite constructions under their stated compatibility and domain assumptions. Related formal tower interfaces are not a Lean certificate for the whole arithmetic model.

Bibliographic scope

The bibliography supplies context for the constructions and records direct comparisons where supported. It is not an exhaustive specialist literature-priority audit. The book does not infer historical priority, independent reproduction, or empirical preference from the extent of the material or from its internal checks.

Source correspondence and reproduction

The accompanying source concordance maps reading passages to the implementation notes moved out of the mathematical narrative. The editorial record preserves the provenance of the source texts, unchanged proof material, document-build results, and known verification limits. Source files, displayed figure assets, and assembly scripts allow the reading volume to be rebuilt. Compiler replay and a full numerical reproduction are separate tasks with their own dependency requirements.

The Foundations Formalization

Exact source correspondence through the first sixteen blocks

1 What this companion makes inspectable

The foundations formalization establishes a precise route from finite projector resolutions to domain-sensitive infinite reconstruction and finite matrix differential calculus. Its first sixteen blocks contain 91 mathematical modules and 1,731 inventoried named declarations. Of these, 1,461 are theorem declarations, 255 are definitions, nine are structures and six are abbreviations. The total includes 47 declarations recovered from Paper VII, including 39 theorems. These are counts of a working mathematical library, including support lemmas and exact examples. They do not count independent physical discoveries.

This companion explains the principal interfaces and how their hypotheses limit a claim. The accompanying `editorial/foundations/lean_crosswalk.csv` lists every one of the 1,731 declarations with its fully qualified name, declaration kind, module, exact source line, source header and file hash. The headers are copied from the source. Section variables and local notation retain their original context, so the CSV is a source crosswalk rather than an independently elaborated collection of `#check` outputs. The original modules remain the authoritative statements.

All source correspondence in this document refers to the foundations files preserved in the Block49 archive. The 91-module foundations audit reported here is a source-identity and receipt audit, not a new compilation of that entire foundations collection. The separately recorded generator replay does not enlarge this audit into a complete foundations replay. The inherited evidence is unusually direct: all 91 foundation module hashes match the retained successful replay of the Block16 frontier at `verification/block17_baseline_full/report.json`. All 114 log files named by that receipt are present and match its recorded hashes. This is a fresh identity check of historical evidence, not a fresh kernel run.

1.1 Three kinds of statement

An abstract interface can assume an operator law, a trajectory or a convergence condition, and then prove consequences of those inputs. A construction supplies an actual object satisfying the interface. An exact example can establish that a hypothesis matters or that a proposed stronger claim fails. All three occur in this library. Their logical roles should remain visible when the same subject is discussed in ordinary prose.

For instance, the finite continuity theorem takes a trajectory whose vector derivative is already known. It proves the derivative of each projected squared norm. It does not assume that projected derivative as an input, but neither does it construct the trajectory. Block11 later constructs global evolution for an arbitrary real diagonal spectrum and proves the precise derivative-domain boundary. That later existence result has a narrower operator class than the earlier conditional balance theorem.

The figures in the foundations atlas illustrate these structures. Their numerical checks have a separate reproduction record. A small numerical residual supports the accuracy of a plotted example. A Lean theorem establishes the quantified statement encoded in its type. Neither record is silently substituted for the other.

2 Types, conventions and the first logical boundary

The finite sector label type is usually called I. Its `Fintype` instance makes sums finite. The ambient complex inner-product space H need not be finite-dimensional. A bounded complex-linear endomorphism belongs to $\mathcal{B}_{\mathbb{C}}(H)$, represented in Lean by a continuous linear map. Completeness is requested

by the constructions that actually use it, including adjoints and square-root dilation. It is not an extra hypothesis of every finite algebraic identity.

The inner product is conjugate-linear on the left and linear on the right. Thus, with $A = -iB$, the receiving-first transfer current is

$$J_{i\leftarrow j} = 2\,\mathrm{Re}\,\langle P_i\psi, AP_j\psi\rangle = 2\,\mathrm{Im}\,\langle P_i\psi, BP_j\psi\rangle\,.$$

The plus sign is part of the convention. The declaration `KomorebiPhysics.hamiltonian_current` proves this algebraic conversion without requiring Hermiticity of B. Hermiticity enters the skewness and conservation conclusions, not this conversion alone.

Formal object	Mathematical meaning
`PacketFamily`	A finite idempotent resolution in an arbitrary ring. This algebraic layer does not itself supply an inner product.
`Atlas H I`	Bounded projectors on a complex inner-product space, with orthogonality, identity resolution and self-adjointness.
`IsSkew A`	The identity $\langle Ax, y\rangle = -\langle x, Ay\rangle$ for all vectors. No unbounded-domain interpretation is hidden here.
`HasDerivAt psi v t`	An actual real-time derivative at t, in the ambient normed space.
`EffectAtlas H I`	Positive Hermitian bounded effects whose finite sum is the identity. Effects may overlap and need not be idempotent.
`StrongAtlas H I`	An orthogonal resolution with `HasSum` over the directed set of finite subsets. Infinite index types are allowed.
`LinearPMap`	A linear operator on an explicit submodule domain. An input to the operator carries its domain-membership proof.
`DomainAtlas T I`	A strong atlas whose individual projections preserve the explicit domain of T.
`GraphCompatible`	Convergence of the finite projector sums in the actual graph norm. It is additional to strong completeness and domain invariance.

A matrix in the finite mixed-state modules is an actual finite complex matrix. Positivity, Hermiticity, normalization and a differential equation remain distinct assumptions. The trace laws are established by matrix algebra. A positivity hypothesis at a time is not, by itself, a positivity-preservation theorem for an evolution.

3 The sixteen-block reading map

The following counts are increments assigned to each block. Block1 includes the Paper VII foundation. The CSV gives the complete module breakdown and the JSON source map records the source identity and evidence boundary.

Block	Main responsibility	Names	Theorems
1	Sharp atlases, transfer, finite refinement	96	76
2	Moving frames and whole-object deformation	64	51
3	Nonlinear and mixed flows, refinement towers	103	83

Block	Main responsibility	Names	Theorems
4	Tensor marginals and information transport	109	94
5	Lindblad jumps and detailed balance	118	98
6	Refinement cancellation and error amplification	89	81
7	Positive effects and moving sources	109	89
8	Parseval channels, dilation and current cycles	120	98
9	Explicit domains and infinite continuity	144	113
10	Bounded spectral surrogates and cutoff topologies	132	113
11	Constructed diagonal evolution and domain boundary	88	81
12	Infinite cells and certified current sums	124	106
13	Residue projectors and unbounded ladders	101	90
14	Domain-correct resolvent commutators	105	87
15	Strong, compact, Hilbert–Schmidt and graph reconstruction	124	111
16	Pointwise no-go and matrix corner differential calculus	105	90
	Total	1,731	1,461

The block order is a useful reading route, but it is not a single chain in which every theorem needs all previous results. Matrix information transport, extrapolation and infinite-domain constructions have their own branches. Imports identify actual dependencies. The module map records their location without promoting chronological order into a mathematical dependence.

4 Blocks 1 and 2: reconstruction, transport and motion

4.1 What the atlas assumes and what it proves

The inherited namespace `KomorebiPaperVII.PacketFamily` supplies the purely algebraic operator corners and their finite reconstruction. Its `reconstruct` theorem retains the operator order

$$X = \sum_i \sum_j P_i X P_j.$$

The receiving index is on the left. The related circuit and path-expansion theorems insert resolutions between consecutive operators. This finite ring calculation does not require norm estimates or a continuum limit.

The structure `KomorebiPhysics.Atlas` extends that resolution by self-adjointness of the projectors. It does not include Parseval as a field. The declarations `Atlas.reconstruct_state`, `Atlas.packet_orthogonal`, `Atlas.inner_reconstruct` and `Atlas.norm_sq_reconstruct` derive vector reconstruction, orthogonality, inner-product reconstruction and

$$\|\psi\|^2 = \sum_i \|P_i\psi\|^2 .$$

Here and below, a name beginning with `Atlas.` is relative to `KomorebiPhysics`.

The theorem `Atlas.hasDerivAt_packet_energy` begins with $\dot{\psi}(t) = A\psi(t)$, expressed as an actual `HasDerivAt` hypothesis. It differentiates the bounded projected vector and then its squared norm. Skewness is not needed for this local derivative calculation. Skewness is needed by `Atlas.current_antisymm`, which allows internal currents to cancel.

For a finite receiving region S, `Atlas.hasDerivAt_subset_energy` proves

$$\frac{d}{dt}\sum_{i\in S}\|P_i\psi(t)\|^2=\sum_{i\in S}\sum_{j\notin S}J_{i\leftarrow j}(t).$$

The derivative on the left is a conclusion. It is derived from the vector equation and finite cancellation. The global statement `KomorebiPhysics.total_energy_constant` assumes the vector equation at every real time and uses the real mean value theorem. An all-time solution remains an input.

4.2 Coarsening is a construction

An address map $r : I \to A$ defines $Q_a = \sum_{r(i)=a} P_i$. Empty fibers cause no difficulty. The definition `Atlas.coarsen` proves the coarse projector laws from the fine ones. The declarations `Atlas.coarsen_corner`, `Atlas.coarsen_energy` and `Atlas.coarsen_current` identify the actual coarse quantities with the appropriate fine sums. Operator corners require a double sum over the receiving and source fibers.

This matters for incidence notation. A matrix of zeros and ones does not automatically encode a partition. An actual address map assigns each fine sector to exactly one coarse label. Duplicating one fine label into two coarse sectors changes the normalization condition and lies outside that constructor.

4.3 Moving coordinates and whole-object rates

Block2 distinguishes unitary frame motion from arbitrary invertible similarity. In `KomorebiPhysics.MovingAtlas`, `unitaryAtlas` constructs an orthogonal transported resolution using the required unitary identities. `hasDerivAt_moving_coordinates` derives the derivative of the actual transformed vector $U(t)^*\psi(t)$. The connection term comes from differentiating U, not from imposing the desired transformed equation.

The namespace `KomorebiPhysics.Deformation` allows nonunitary invertible transport. It preserves an algebraic idempotent resolution, but it does not claim that every transported idempotent stays self-adjoint. Its whole-object theorem differentiates

$$P(t)X(t)Q(t)$$

with all three product-rule terms. The declarations `hasDerivAt_corner` and `hasDerivAt_whole_object` retain intrinsic motion of X and the two moving frame factors. The inverse-curve interface `hasDerivAt_ring_inverse` proves the inverse derivative under completeness and invertibility at the base point.

The exact shear examples make these distinctions visible. In the example with corner contributions $4N, -3N, 2N$, the full rate is $3N$. Deleting any contribution changes the answer. A co-moving object can have a nonzero laboratory derivative, and a nonunitary similarity can produce an idempotent that is not symmetric. Those are exact finite examples, not floating-point diagnostics.

5 Blocks 3 to 5: mixed dynamics and information

5.1 State dependence does not erase the input equation

The namespace `KomorebiPhysics.NonlinearFlow` replaces a fixed bounded generator by a state- and time-dependent one. The packet derivative is still derived from an actual trajectory equation. Skewness is required at the trajectory point for current antisymmetry and total conservation. The declaration `total_constant` therefore describes conservation along a supplied global solution. It is not an existence theorem for arbitrary nonlinear equations.

In `KomorebiPhysics.MixedTransfer`, the current and probability are finite matrix expressions. The real-valued probability is identified with the literal complex trace under the stated Hermiticity assumptions. The declaration `MatrixAtlas.hasDerivAt_probability` starts from the derivative of the density matrix, not from a postulated scalar probability equation. The exact rank-one specialization `MatrixAtlas.pure State_current` connects the mixed-state current to the vector current.

5.2 Tensor marginals and entropy

Block4 constructs the tensor atlas from actual local finite matrix atlases. The marginal formulas hold for a general joint matrix. Product factorization requires a product input and is stated separately. The `KomorebiPhysics.PartialTrace` constructions prove positivity and trace preservation of the reduced matrices under the appropriate joint hypotheses, and identify the resulting local probabilities with the joint marginal sums.

The information functionals in `KomorebiPhysics.InformationTransport` are functionals of packet probabilities. The entropy is the Shannon expression

$$S(p) = -\sum_i p_i \log p_i,$$

using the natural logarithm. It is not the von Neumann entropy of an arbitrary density operator. The derivative formulas require strict positivity at the base time. Their use of a logarithm is not justified merely because Lean gives total functions to division and logarithm.

For an antisymmetric current satisfying the component balance law, the entropy derivative is the symmetrized current expression proved by `hasDerivAt_entropy_current`. This identity alone does not give a sign. The fixed real exponent in the power-sum and Rényi statements is not restricted to integers. The Rényi theorem also requires a nonempty index type and $q \neq 1$. The limit as $q \to 1$ is a separate statement.

The finite entanglement witness has similarly explicit scope. The Bell example excludes every explicit finite nonnegative mixture of pure product states through a negative partial-transpose pairing. It is an exact witness within the represented finite matrix setting. Its theorem name is `KomorebiPhysics.QuantumInformationExamples.bell_not_finite_product_mixture`.

5.3 What detailed balance adds

Block5 constructs jump matrices with square-root rate amplitudes, then computes the gain and anti-commutator terms on arbitrary coherent matrix inputs. The scalar population equation follows from the original matrix derivative. Off-diagonal coherences have the half-sum of escape rates prescribed by the actual jump convention. These statements do not supply an integrated positive solution of every Lindblad equation.

In `KomorebiPhysics.DetailedBalance`, the fixed positive reference π and the detailed-balance identity are explicit. Along a positive supplied population trajectory, the relative entropy

$$D(p\|\pi) = \sum_i p_i \log \frac{p_i}{\pi_i}$$

has a nonpositive derivative. The declarations `dissipation_nonneg`, `relativeEntropy_deriv_nonpos` and `relativeEntropy_antitoneOn` distinguish the algebraic sign, the pointwise derivative and monotonicity on a convex time set.

The conductance form covers inactive edges. Its flux-ratio presentation uses an explicit convention there and proves equivalence on active edges. It does not rely on evaluating an undefined ordinary ratio of two zero fluxes. The equality case first states equality of ratios on active edges. Identifying the whole distribution with the reference requires positivity of every distinct-label rate and normalization of both distributions, as required by the corresponding theorem. Arbitrary antisymmetric transfer does not imply this entropy conclusion.

6 Blocks 6 to 8: cancellation and positive measurements

6.1 Finite-order cancellation with its actual error budget

The namespace `KomorebiPhysics.RefinementCancellation` defines the sampled table and its extrapolation recurrence. It accepts a scale ratio $b > 1$. The factors cancelling the first terms are proved algebraically, and all recurrence denominators are proved nonzero. The analytic statements in `RefinementAsymptotics` then combine a finite expansion with a genuine asymptotic remainder.

For a fixed finite cancellation depth, `dyadic_order_improvement` and `positive_step_order_improvement` improve the order under their input expansion hypotheses. The one-sided statement does not require the original expansion to extend to negative step sizes. Neither theorem proves convergence of an infinite formal series simply because the source displays one.

The bounds in `RefinementBounds` retain amplification of sample error. The exact example `KomorebiPhysics.RefinementExamples.sampleNoise_amplified` shows why cancellation cannot be advertised as a universal contraction of noisy data. `RefinementAsymptotics.tower_not_extra_order` gives a genuine asymptotic sharpness statement beyond the cancelled degree. Improving the truncation order and controlling perturbations are separate obligations.

6.2 Effects, expectations and moving sources

In Block7, `KomorebiPhysics.EffectAtlas` replaces orthogonal projectors by positive Hermitian effects E_i with $\sum_i E_i = I$. Its probability is

$$p_i = \langle \psi, E_i \psi \rangle ,$$

interpreted as a real number by a proved reality statement. It generally differs from $\|E_i\psi\|^2$. Positivity and normalization give the probability bounds. They do not imply sharpness.

The effect current is explicitly antisymmetrized. Its antisymmetry therefore holds by construction, while identifying its row sum with the actual state derivative uses the skewness condition. If the effects move, `EffectAtlas.hasDerivAt_moving_probability` retains an additional source

$$s_i = \left\langle \psi, \dot{E}_i \psi \right\rangle .$$

Differentiating the identity resolution proves that the sources sum to zero. Omitting them would change the moving-measurement law even when total probability remains constant.

A fixed stochastic postprocessing matrix also has an exact constructor. Its weights must be real, nonnegative and column normalized. The resulting effects and probabilities obey the weighted aggregation formulas, and the current matrix becomes RJR^{T}. Time-dependent weights would contribute further derivative terms and are not hidden in this fixed-weight theorem.

6.3 Dilation is universal within the finite bounded setting

Block8 begins with actual bounded channels $C_i : H \to K_i$ satisfying a Parseval resolution. The namespace `KomorebiPhysics.ParsevalFamily` proves state and operator reconstruction, constructs analysis and synthesis maps, and proves that synthesis is the adjoint of analysis. The finite direct sum uses its Hilbert norm.

The square-root construction then supplies channels for an arbitrary finite effect atlas on a complete complex Hilbert space. The declarations `EffectAtlas.sqrtFamily` and `EffectAtlas.exists_finite_dilation` close the factorization obligation at this level. Finite-dimensionality of H is not required. Countably many outcomes and unbounded effects require additional work.

Write the analysis isometry as V. The lifted identity is the projection $Q = VV^*$ onto its image. The supported lift $X \mapsto VXV^*$ is multiplicative, but general compression $Y \mapsto V^*YV$ need not be. The exact compression-defect theorems explain the difference. In particular, probability reconstruction from all channel corners cannot be replaced by retaining only diagonal corners of an arbitrary operator.

Finally, `KomorebiPhysics.CycleGauge.difference_cycle_iff` describes the ambiguity of a current determined only by continuity. Among antisymmetric arrays, two currents have equal divergence exactly when their difference has zero row sums. The triangle example gives a nonzero such cycle. The state balance law fixes divergence and boundary transfer, while leaving this cycle freedom.

7 Blocks 9 and 11: domains and actual diagonal evolution

7.1 Three assumptions that must not be merged

The infinite interface separates strong completeness, domain invariance and graph compatibility. A `StrongAtlas` supplies genuine `HasSum` statements for projected vectors. A `DomainAtlas` also proves that each projector sends a vector in $\mathcal{D}(T)$ back into $\mathcal{D}(T)$. Its `GraphCompatible` predicate additionally demands convergence of finite sums in

$$\|x\|_T^2 = \|x\|^2 + \|Tx\|^2 .$$

The graph norm is constructed from the Hilbert product norm. It is not the default ambient norm on the domain subtype.

The exact equivalence `KomorebiPhysics.GraphDomain.graph_convergence_iff` identifies graph convergence with convergence of both the vectors and their operator images. Strong convergence of the projectors alone cannot justify applying an unbounded operator to a limit. The theorem `DomainAtlas.graph_iff_action` isolates the missing action-sum condition explicitly.

The domain current theorems assume formal symmetry of T, which is weaker than self-adjointness. The fixed-row theorem `DomainAtlas.current_hasSum` proves a genuine scalar current sum from symmetry, domain preservation and strong reconstruction. It does not assert joint summability of the two-

index array. The graph-compatible formulation supplies a separate full-action limit and leads to actual finite-region boundary derivatives.

7.2 Constructing the maximal multiplier

The namespace `KomorebiPhysics.DiagonalDomain` constructs the real diagonal operator with eigenvalue sequence λ_n. Its maximal domain consists of those $x \in \ell^2$ for which $(\lambda_n x_n)_n \in \ell^2$. Self-adjointness and closedness are proved for this concrete partial operator. The domain is not declared to be all of ℓ^2, and a formal symmetry equation alone is not relabelled as self-adjointness.

The spectral projector atlas supplies an actual graph-compatible realization. Weighted tail series control convergence of the vector and its generator image. The construction is a spectral model. Identifying it with a chosen differential oscillator, together with that operator's domain and boundary conditions, is a different theorem.

7.3 The point at which existence is supplied

The modules `SpectralTimeContinuity.lean`, `SpectralSchrodinger.lean`, `SpectralSolutions.lean` and `SpectralTimeApproximation.lean` share the namespace `KomorebiPhysics.SpectralEvolution`. Filenames alone therefore do not determine the namespace of the theorem to inspect.

For the real diagonal spectrum, the orbit is coordinatewise phase multiplication,

$$(U(t)x)_n = e^{-it\lambda_n} x_n.$$

The theorem `continuous_orbit` proves strong time continuity for every ℓ^2 initial state. `domain_preserved` proves invariance of the maximal domain. `orbit_hasDerivAt` proves the actual ambient derivative on that domain. The declaration `exists_unique_solution` gives global domain solutions, while `differentiableAt_iff_mem_domain` establishes the exact boundary:

$$t \longmapsto U(t)x \text{ is differentiable at } t_0 \quad \Longleftrightarrow \quad x \in \mathcal{D}(T).$$

The point t_0 may be any real time. Norm preservation is available for every initial state, but differentiability is not. This is why rough-state conservation cannot be justified by differentiating a nonexistent ambient derivative.

The cutoff estimates `cutoff_uniform_in_time` and `cutoff_velocity_uniform` retain one fixed initial state. The velocity estimate requires domain membership. They do not give a uniform finite-cutoff approximation over the entire unit ball. The exact rough-state and cutoff-velocity examples verify that these distinctions have content.

8 Block 10: which topology a cutoff controls

The bounded diagonal module builds an operator from a bounded scalar symbol and proves its supremum norm. The truncation module then converts operator questions into exact scalar tail questions. The harmonic spectral examples instantiate this machinery, including a resolvent at real parameter zero because that particular model has a spectral gap.

Construction	Checked behavior	Boundary retained
Harmonic resolvent cutoff	Operator-norm tail tends to zero, with an exact endpoint formula.	Domain range and both inverse identities are proved.
Heat multiplier at fixed $t > 0$	Operator-norm cutoff convergence.	At $t = 0$, cutoff error is one. Negative time has an unbounded symbol.
Phase multiplier	Strong cutoff convergence for each state.	Operator-norm cutoff error is one at every finite cutoff.
Anchored bounded transform	Norm convergence with the correct scalar tail anchor.	Zero padding has error one for the source transform.
Anchored Cayley transform	Norm convergence with the correct anchor and exact tail formula.	A wrong anchor leaves a nonvanishing error.

The generic criterion is `KomorebiPhysics.SpectralTruncation.anchored_norm_iff`. It states that the anchored truncations converge in norm exactly when the bounded symbol tends to that anchor. In mathematical notation, the approximation is

$$S_M D_a S_M + c(I - S_M).$$

The exterior value c is part of the operator. It cannot be dropped just because the retained finite block has the desired entries.

These are exact infinite-dimensional statements. They are stronger than checking a few large matrices and more discriminating than the unqualified phrase "the truncations converge." They still concern the constructed diagonal realization, rather than a general functional calculus for every self-adjoint partial operator.

9 Blocks 12 and 13: infinite cells and residue routing

9.1 The existence of the sum is part of the result

In Block12, `StrongAtlas.cell` forms a coarse cell from an arbitrary fiber, including an empty or infinite fiber. `StrongAtlas.coarsenInfinite` constructs the coarse strong atlas. Positivity, self-adjointness, idempotence and boundedness of the coarse projectors are proved from the fine resolution.

Preserving an unbounded domain needs more than ambient convergence. For a closed operator and a graph-compatible fine atlas, `DomainAtlas.cell_mem_domain` proves membership of the coarse projected vector. The alternative constructor `coarsenOfPreserves` accepts an independently established preservation result. The source does not infer closure of a domain under arbitrary ambient limits.

For coarse currents there are two distinct formulations. The iterated theorems prove both inner and outer sums in the specified order. The product-index theorems require joint summability of the corresponding scalar array. The declaration `KomorebiPhysics.CurrentSummationExamples.not_joint_summable` supplies an abstract antisymmetric counterexample with zero iterated sums in both orders but no joint summability. It is not presented as a particular Hamiltonian realization.

The actual coarse derivative is obtained by differentiating a bounded coarse projector applied to the original trajectory. It is not inferred by an unproved exchange of an infinite sum and a derivative. In the diagonal constructed-solution specialization, the earlier existence theorem supplies the trajectory equation needed here.

9.2 Residues are bounded projectors, ladders are partial operators

Block13 uses a positive modulus through a `NeZero` assumption. `KomorebiPhysics.ResidueProjectors.projector_number_average` identifies the finite character average with the residue projector, and `projector_eq_mask` identifies it with the infinite coordinate mask. The normalization factor and the sign of the receiving character are retained. The infinite subseries has a proved `HasSum` statement.

The ladder operators require their own domain. The names `KomorebiPhysics.LadderDomains.lower` and `KomorebiPhysics.LadderDomains.raise` designate actual partial operators with square-root amplitudes. Their common domain is the first-moment condition described by `domain_iff`. The adjoint theorems establish genuine unbounded adjoint equalities, and closedness is a consequence proved for these operators.

Routing then states that lowering moves a residue label down by one and raising moves it up by one, with labels interpreted in $\mathbb{Z}/M\mathbb{Z}$. The input remains domain-valued. Vacuum behavior is handled explicitly. A second ladder application would require membership in the relevant product domain, so a canonical commutator theorem is not inferred from one-step routing alone. Likewise, coherent-state factorial weights and a differential oscillator identification are beyond the statements represented here.

10 Blocks 14 and 15: resolvents and reconstruction

10.1 A certificate keeps the unbounded domain visible

For a partial operator T, a bounded map R is accepted as a resolvent only with domain range and both inverse identities for $T - zI$. A bounded commutator extension B agrees with $[T, P]$ on the actual domain, and P preserves that domain. Under those conditions, `KomorebiPhysics.ResolventDomain.Certificate.commutator_identity` proves

$$[R, P] = -RBR.$$

This identity does not apply T off its domain. The bounded map P need not be a projection and need not reduce T. Its operator-norm consequence is

$$\|[R, P]\| \leq \|R\|^2 \|B\|.$$

For a symmetric operator and a genuine nonreal resolvent, the additional estimate $\|R\| \leq |\operatorname{Im} z|^{-1}$ is derived. Nonzero imaginary part is explicit.

The namespace `KomorebiPhysics.DiagonalResolvents` constructs the certificates for arbitrary real diagonal spectra and nonreal z. Generic certificate consequences therefore have a concrete realization, but they do not assert resolvent existence for every merely symmetric partial operator.

The converse theorem `Certificate.commutator_zero_iff` concerns domain invariance and vanishing of the generator commutator on that domain. The exact nonreducing-packet examples show nonzero leakage. Formal symmetry, reducing structure and a small resolvent commutator remain different properties.

10.2 Strong reconstruction and its upgrades

For a bounded X, uniformly contractive cutoffs S_α converging strongly to the identity give strong convergence of $S_\alpha X S_\alpha$ to X. The theorem `KomorebiPhysics.StrongReconstruction.compression_strong` states this at the level of each fixed vector.

Compactness of X permits an operator-norm upgrade under the self-adjoint contraction hypotheses used in `KomorebiPhysics.CompactReconstruction.compact_compression`. Self-adjointness is a sufficient hypothesis of this theorem. The library does not claim it is necessary in every possible formulation. The proof does not obtain adjoint convergence for an arbitrary strongly convergent sequence by assumption-free inference.

The Hilbert–Schmidt layer begins by proving summability of actual bounded-operator columns. The predicate `KomorebiPhysics.HilbertSchmidtCore.IsHS` is necessary when interpreting the scalar formula `hsNorm` as the norm under discussion. Basis independence, adjoint invariance and ideal bounds are proved. This is substantive Hilbert–Schmidt infrastructure, but it is not advertised as a complete bundled theory of all Schatten classes.

The exact geometric diagonal example has squared Hilbert–Schmidt norm $4/3$, residual squared Hilbert–Schmidt norm $(4/3)4^{-M}$, and residual operator norm 2^{-M}. The errors are nonzero at every finite cutoff. Exact tail formulas and actual limits are represented independently of plots.

10.3 Graph reconstruction does not iterate the operator

The declaration `KomorebiPhysics.GraphReconstruction.reconstructed` uses a domain-valued finite cutoff before applying T, then applies the ambient cutoff to the result. Under graph compatibility,

$$S_\alpha T S_\alpha x \longrightarrow Tx \qquad (x \in \mathcal{D}(T))$$

in the ambient norm. This does not assert that Tx belongs to $\mathcal{D}(T)$, or that these outputs converge in the graph norm of T. Either strengthening would require more domain control.

The actual atlas cutoffs are orthogonal projections and contractions, but a finite number of retained sectors does not imply finite-dimensional ranges. The library also contains a compact rank-one counterexample to dropping the needed cutoff hypotheses from an arbitrary sandwich theorem. This prevents a true strong statement from being silently promoted to an operator-norm statement.

11 Block 16: the finite differential boundary

The pointwise algebra in `KomorebiPhysics.PointwiseNoGo` is $\mathbb{C}^I$ for a finite label type, with ordinary coordinatewise multiplication. A derivation is an actual complex-linear map obeying Leibniz. The declarations `atom_derivative_coordinate`, `map_atom_eq_zero` and `derivation_eq_zero` prove that it must vanish.

The proof tests Leibniz on each coordinate idempotent. Each coordinate of its derivative is forced to be zero. Finite reconstruction by those idempotents then kills the derivative of every element. The restriction to a finite product of copies of $\mathbb{C}$ matters. It is not a theorem that every finite-dimensional commutative algebra has only the zero derivation.

The matrix alternative is explicit. In `KomorebiPhysics.CornerDerivation`,

$$\delta_K(X) = i(KX - XK)$$

is defined by matrix operations. The theorem `delta_mul` proves Leibniz. Hermitian K gives star compatibility and preservation of Hermitian inputs. Trace cyclicity gives

$$\mathrm{Tr}((\delta_K X)Y) = -\,\mathrm{Tr}(X\delta_K Y).$$

The conclusion is a finite matrix integration-by-parts identity. No spatial integral or continuum boundary condition has been inserted implicitly.

With $F_{KL} = i[K, L]$, the namespace `KomorebiPhysics.CornerCurvature` proves

$$[\delta_K, \delta_L]X = \delta_{F_{KL}} X.$$

The Jacobi identity controls the factor of i and its sign. The checked central-curvature implication distinguishes the matrix value from its induced inner action. Exact two-by-two examples establish a nonzero derivation and a nonzero curvature action.

The namespace `KomorebiPhysics.CornerReconstruction` then reconstructs the derivative from projector corners. Inserting an intermediate resolution into both commutator terms gives the path expression. Coarse corners are exact sums of the corresponding fine corners. This is the foundations endpoint: a precise finite noncommutative differential calculus with reconstruction. Its derivations do not by themselves construct continuum spacetime, interacting gauge evolution or gravity.

12 Reading and replaying the evidence

12.1 Use the exact receipt, not a generic filename

The historical Block16 proof guide points to `verification/final/report.json`. In the Block49 package that filename now holds a later cumulative run. The retained `verification/block17_baseline_full/report.json` instead has exactly the Block16 totals and hashes. It records a successful run on 14 September 2026, an axiom audit of 1,731 declarations and 115 successful negative controls. The allowed axiom dependencies are `propext`, `Classical.choice` and `Quot.sound`.

The receipt's stated trust boundary is fresh local proof terms checked against the restored pinned cache, with imported Lean and Mathlib libraries not independently rebuilt in that run. For the accompanying record, all 91 mathematical source hashes and all 114 referenced log hashes were compared with the retained receipt and matched. A read-only, comment- and string-masked scan of the same mathematical files found none of the listed admission, custom-axiom, unsafe/native-evaluation or selected bypass tokens. Such a scan complements source identification. It cannot replace kernel elaboration.

The negative controls are deliberately invalid examples kept apart from mathematical modules. Their success condition is an expected rejection with the appropriate diagnostic, rather than a proof of a false statement. Counts of those controls describe tests of the checking procedure and hypothesis boundaries. They do not add 115 mathematical results to the theorem total.

12.2 A compact inspection recipe

The pinned toolchain is Lean `v4.34.0-rc2`. The package manifest pins Mathlib to the revision beginning `aeb72f1e9981`. A reader can inspect selected declarations without importing the entire later physics frontier. For example, an inspection file can contain:

```
import KomorebiPhysics.Transfer
import KomorebiPhysics.SpectralSolutions
import KomorebiPhysics.CornerCurvature

#check KomorebiPhysics.Atlas.hasDerivAt_subset_energy
#check KomorebiPhysics.SpectralEvolution.exists_unique_solution
#check KomorebiPhysics.CornerCurvature.derivation_commutator

#print axioms KomorebiPhysics.Atlas.hasDerivAt_subset_energy
#print axioms KomorebiPhysics.SpectralEvolution.exists_unique_solution
#print axioms KomorebiPhysics.CornerCurvature.derivation_commutator
```

This is a reading and replay recipe, not a record of a command run during the accompanying record. A replay should use the pinned environment and keep its own logs. The original source archive includes the historical guides, audit machinery, negative controls and dependency manifest. It does not make the present PDF a substitute for that environment.

For the full physics frontier, the root `KomorebiPhysics.lean` aggregator in the Block49 archive stops at Block47 and omits the two latest modules. Some generic wrapper metadata also retains earlier versions. This does not alter the 91 foundation files whose identities were compared here, but it prevents a root import alone from serving as evidence that the entire Block49 frontier was freshly replayed. The separate Dirac Lean companion gives the exact later module names.

12.3 How to use the machine crosswalk

The CSV's `fully_qualified_name` column is the lookup key. Its `module` column identifies the import, while `source_file` and `source_line` identify the original declaration. These are intentionally separate because several modules share one namespace. `source_header` retains the literal declaration header and `source_sha256` identifies its containing file. The axiom column is copied from the inherited matching receipt.

The JSON source map records the block counts, all module identities, receipt identity, log comparisons and scan scope. It also records that no fresh Lean check was run. The extraction script is a read-only reproduction aid. Its source-header output is not an elaborator, and its 1,731-name boundary follows the existing named-declaration inventory rather than counting every generated structure projection or compiler-generated constant.

To assess a manuscript sentence, locate the corresponding declaration, read its surrounding section variables, then distinguish the hypotheses from the conclusion. For an infinite statement, check the summability or filter topology. For a partial operator, check the domain subtype. For an evolution, determine whether the trajectory is supplied or constructed. These questions turn the formal companion into a usable specification of the mathematics rather than a list of theorem names.

Formal Verification of the Free Dirac Operator

A source guide to the checked construction and its proof boundaries

1 What this companion certifies

This companion identifies the checked free Dirac Hamiltonian, its strong generator, and its centered approximation on the flat unit three-torus. It records the declarations, their hypotheses, and the boundary between written proofs and formal verification. The operator has an exact first-order Fourier domain, self-adjointness, closedness, and an energy identity. The generator identifies that full domain with strong differentiability. The centered extension supplies graph and resolvent limits and quadratic bounds under a third-order Fourier moment.

The source release is classified as **checked modular cumulative**. That classification has a precise meaning. The new Block 49 module and its import closure have recorded successful checks, while earlier source files are retained under preservation records. It does not mean that every historical manuscript sentence has become a theorem, or that a single aggregate command has freshly checked every file in this editorial edition.

The historical verification tables below describe the preserved v0.49.0 handoff. The preceding generator checkpoint rebuilt the exact operator target and checked the generator extension against the same pinned environment. The sections below distinguish that checkpoint from the current centered-extension record. The inherited Lean source files remain unchanged. Historical receipts are not silently relabelled as current runs.

The mathematical source, its verification records, and its explanatory wrappers are different objects. Some generic baseline wrappers describe earlier checkpoints. The named source targets and their individual receipts govern the operator and generator claims. Package labels and aggregate inventory counts do not replace those records.

1.1 A reading convention

The words *definition*, *hypothesis*, and *theorem* have separate roles here. A definition introduces an object, often together with proof fields required by a Lean structure. A hypothesis records what a theorem needs from its inputs. A theorem derives a conclusion from those inputs. For example, the early atlas laws are assumptions defining an admissible atlas. Reconstruction and coarse balance are consequences of those laws. A supplied evolution equation is not itself an existence theorem for that evolution.

Names printed below are declaration names found in the supplied source. A short name is relative to the namespace explicitly stated for its section. File paths are relative to the root of the checked source archive. The machine-readable companion, `editorial/lean_crosswalk.csv`, includes the full namespace, declaration kind, exact source path, line, statement signature, available declaration comment, and source hash for every declaration extracted from Blocks 41–49. It contains 322 named declarations, of which 207 are theorems or lemmas. These are source-extraction counts for those nine modules, not independent counts of physical laws.

2 The mathematical objects and their Lean types

The final construction uses a fixed flat geometry and a fixed four-component spinor representation. The spatial torus is the unit torus, so its integer Fourier mode $n \in \mathbb{Z}^3$ has momentum $k_n = 2\pi n$. The mass parameter is any real number. No positive-mass restriction is required by the final domain theorem.

The implementation separates the algebraic four-spinor representation from the Hilbert-space representation. This is necessary because a function type with four complex coordinates does not, by itself,

assert the Euclidean norm needed by the energy and unitarity arguments. The conversion maps preserve the coordinate values, while the Hilbert type supplies the correct norm and inner product.

Mathematical object	Source representation and interpretation
Fourier label	`Mode:=Fin3` $\to \mathbb{Z}$. The set of labels is infinite.
Algebraic spinor	`Fin4` $\to \mathbb{C}$. Used for explicit Clifford formulas.
Hilbert spinor	`EuclideanSpace` $\mathbb{C}$ `(Fin4)`. The norm is the Euclidean four-spinor norm.
Scalar torus field	`Lp` $\mathbb{C}$ `2(volume:Measure(UnitAddTorus(Fin3)))`. Equality is equality of L^2 classes.
Physical spinor field	`PiLp2(fun_:Fin4=>ScalarTorusL2)`. This is the four-coordinate Hilbert sum used for $L^2(\mathbb{T}^3;\mathbb{C}^4)$.
Coefficient field	`CoeffL2`, the space $\ell^2(\text{Mode};\text{HilbertSpinor})$. Square summability is already part of the type.
Bounded mode Hamiltonian	A continuous complex-linear endomorphism of one Hilbert spinor space. Each fixed mode has a bounded matrix Hamiltonian.
Full continuum Hamiltonian	A `LinearPMap`, written CoeffL2 $\to_{\mathbb{C}}^{\text{partial}}$ CoeffL2 in mathematical notation or its physical-space counterpart. Its domain is a submodule, and an argument carries a proof of domain membership.
Fourier equivalence	`torusFourierEquiv`, a complex linear isometric equivalence from physical torus fields to coefficient fields.

The inner product convention is conjugate-linear in its first argument and linear in its second. The source uses `inner` $\mathbb{C}$. In particular, the time evolution of a mode is built from the bounded exponential of $-itH_m(n)$. Maintaining this convention is necessary when reading the current signs and the adjoint calculations.

2.1 What the spacing parameter means here

In the Fourier evolution modules, the centered multiplier is

$$k_{h,j}(n) = \frac{\sin(k_{n,j}h)}{h},$$

with the removable value $k_{n,j}$ inserted at $h = 0$. Both the centered and continuum propagators act on the same infinite label space $\mathbb{Z}^3$. Those original limit theorems take $h \to 0$ through positive real values; the centered extension below also supplies the two-sided punctured limit.

This model does not replace $\mathbb{Z}^3$ by the Fourier labels of a finite sampled grid. There is no changing Nyquist cutoff in the final theorem statement, and the proof is not an identification of a finite lattice state space with all of torus L^2. Earlier finite-resolution and geometric refinement constructions are substantial results of their own. A further compatibility theorem is needed before those constructions can be substituted for the centered multiplier family in a full operator-limit claim.

3 The dependency chain from symbols to physical fields

Each of the nine recent modules introduces a distinct part of the argument. Their principal dependency chain follows the block order below. Some modules also import mathematical infrastructure and earlier Clifford constructions. Thus the table is a reading route, not a complete graph of all imported library files.

Block	Module below `KomorebiPhysics/`	Main responsibility
41	`Relativistic3DFourierGenerator.lean`	Centered symbol, exact mass shell, fixed-mode generator consistency.
42	`Relativistic3DFourierL2Generator.lean`	Orthogonal torus embeddings and finite-packet L^2 consistency.
43	`Relativistic3DFourierPropagator.lean`	Bounded mode exponentials and fixed-time packet convergence.
44	`Relativistic3DFourierHilbertStability.lean`	Hilbert self-adjoint mode operators and exact unitarity.
45	`Relativistic3DFourierCompactTime.lean`	Uniform bounds on a fixed compact time interval.
46	`Relativistic3DFullFourierEvolution.lean`	Full coefficient evolution and removal of finite support.
47	`Relativistic3DTorusFourierEvolution.lean`	Surjective Fourier equivalence and physical-space transport.
48	`Relativistic3DUnboundedDirac.lean`	Maximal coefficient domain, adjoint, and graph approximation.
49	`Relativistic3DWeightedDiracDomain.lean`	Exact weighted domain and physical unbounded Hamiltonian.

The namespace usually matches the filename without its suffix. There is one important exception: Block 49 uses

```
KomorebiPhysics.Relativistic3DWeightedDomain
```

although its filename contains `WeightedDiracDomain`. The crosswalk preserves this distinction. Guessing a fully qualified name from the filename would produce the wrong namespace.

3.1 Blocks 41 and 42: exact algebra before limits

In `KomorebiPhysics.Relativistic3DFourierGenerator`, the declaration `centralDifferenceMultiplier_eq` reduces the centered translation multiplier to the sine quotient. The Hamiltonian symbols are defined by `continuumDiracSymbol` and `centralDiracSymbol`. Their square identities are `continuumDiracSymbol_sq` and `centralDiracSymbol_sq`.

For the continuum symbol, the identity is

$$H_m(n)^2 = (m^2 + |k_n|^2)I.$$

The centered symbol has the same formula with each momentum component replaced by its sine quotient. These are exact identities for the chosen Clifford representation. They are not numerical tests of a dispersion curve.

The small-spacing theorem `centralModeAction_sub_continuum_isBigO` fixes the mass, mode, and spinor before taking the limit. Its constant may therefore depend on those fixed inputs. The finite-packet theorem also fixes its finite support. This order of quantifiers prevents a finite-packet estimate from being read as a bound uniform over all Fourier modes.

Block 42 moves these statements into actual torus L^2. In namespace `KomorebiPhysics.Relativistic 3DFourierL2Generator`, `modeEmbedIsometry` embeds a single Hilbert spinor at one Fourier mode. The declaration `inner_modeEmbed` supplies mode orthogonality, and `fourierPacketL2_norm_sq` gives the exact finite Parseval identity. The convergence theorem is `tendsto_centralPacketL2`. Its quadratic precursor is `packetResidualL2_isBigO`.

This step turns coordinatewise symbol control into control of a field norm. It still retains fixed finite Fourier support. There is no unbounded operator acting on arbitrary torus fields at this stage of the chain.

3.2 Blocks 43 and 44: bounded exponentials and stability

In namespace `KomorebiPhysics.Relativistic3DFourierPropagator`, the mode Hamiltonians become continuous linear maps. The declarations `continuumModePropagator` and `centralModePropagator` use the Banach-algebra exponential. The theorem `centralModePropagator_sub_isBigO` gives quadratic convergence at a fixed time. The declaration `tendsto_centralPropagatedPacketL2` transfers this to a fixed finite packet.

The group identities use the fact that two time-scaled copies of the same generator commute. For example, `continuumModePropagator_add_time` proves addition of times, and `continuumModePropagator_ neg_mul` gives the inverse-time identity. Invertibility alone does not supply the Hilbert stability needed to remove Fourier tails.

Block 44 supplies that stability in namespace `KomorebiPhysics.Relativistic3DFourierHilbertSt ability`. The declarations `continuumHilbertModeOperator_selfAdjoint` and `centralHilbertModeOp erator_selfAdjoint` use the Euclidean Hilbert representation. Their exponentials are proved unitary by `continuumHilbertModePropagator_mem_unitary` and `centralHilbertModePropagator_mem_unitary`. The corresponding norm and inner-product preservation theorems apply to every spinor at each fixed mode.

The finite-packet norm identities are `continuumHilbertPropagatedPacketL2_norm_sq` and `centralHilbe rtPropagatedPacketL2_norm_sq`. Stability is exact. It does not acquire an error depending on the number of modes or the time interval.

4 Compact-time convergence and removal of the cutoff

Block 45 changes the time quantifier. In namespace `KomorebiPhysics.Relativistic3DFourierCompactTime`, the theorem `exists_compactTime_mode_propagator_bound` takes m, T, n, with $T \geq 0$, and proves that some $K \geq 0$ works eventually as $h \to 0^+$, simultaneously for every $|t| \leq T$:

$$\left\| U_{m,h}(t,n) - U_m(t,n) \right\| \leq K h^2.$$

The norm here is the bounded operator norm of the mode endomorphisms. The coefficient K may depend on the chosen mode, mass, and time window.

The packet version, `exists_compactTime_packet_propagator_bound`, takes a finite set of modes and a fixed amplitude function. Its constant can also depend on these data. The proof first controls the

exponential on one closed operator ball and then makes the finitely many eventual estimates hold together. The result is an existence bound. The guide does not replace it by a sharp Duhamel constant or an explicitly optimized constant.

The reusable stability lemma is `norm_sub_le_two_tail`. For norm-preserving linear evolutions and any finite approximation g, it gives

$$\|U_h(t)f - U(t)f\| \le 2\|f-g\| + \|U_h(t)g - U(t)g\|.$$

The two tail terms have the same size because both evolutions preserve norm. This coefficient two is carried into the cutoff-removal proof.

4.1 Full coefficient evolution

In namespace `KomorebiPhysics.Relativistic3DFullFourierEvolution`, `modewiseMap` first packages coordinatewise action on square-summable coefficients. The proof that the output remains square summable uses the modewise norm equality. The definitions `continuumCoeffPropagator` and `centralCoeffPropagator` are linear isometries on the full coefficient space. Surjectivity follows by applying the opposite-time operator, yielding the linear isometric equivalences `continuumCoeffUnitary` and `centralCoeffUnitary`.

The finite approximation is `coeffTruncationSf`. It retains exactly the coefficients in a finite set S. The statement `exists_coeffTruncation_approx` says that every square-summable coefficient field has such an approximation with arbitrarily small norm error. No extra regularity is required for that statement.

The final theorem `fullCoeffPropagator_compactTime_converges` has the following mathematical quantifier structure:

$$\begin{aligned} &m \in \mathbb{R}, \quad T \ge 0, \quad f \in \ell^2(\mathbb{Z}^3;\mathbb{C}^4),\\ &\forall \varepsilon > 0, \quad \text{eventually as } h \to 0^+,\\ &\qquad \forall |t| \le T, \qquad \left\|U_{m,h}(t)f - U_m(t)f\right\| < \varepsilon. \end{aligned}$$

The source expresses "eventually" with $\forall^{\mathrm{f}} h$ in $\mathcal{N}[>](0 : \mathbb{R})$. It is strong convergence for each fixed state, uniform over the selected compact time interval. It is not convergence in operator norm over the unit ball of the full coefficient space.

The order of choices explains the rate boundary. First choose the finite support to make the tail small. Next choose the spacing to control that finite core. The finite-core constant may grow when the support grows. Consequently, the theorem proves convergence for every ℓ^2 state without supplying a universal quadratic rate for arbitrary states.

4.2 Physical-space transport

Block 47 proves that the Fourier map is onto. In namespace `KomorebiPhysics.Relativistic3DTorusFourierEvolution`, the declarations `modeEmbed_torusCoeffL2`, `modeEmbed_linearIsometry_surjective`, and `modeEmbed_isHilbertSum` establish the synthesis and surjectivity needed for `torusFourierEquiv`.

This is more than an isometric embedding. An embedding alone would leave a possible missing orthogonal sector in physical space. Surjectivity permits every torus L^2 spinor to be transported to coefficient space and back.

The physical evolutions are `continuumTorusUnitary` and `centralTorusUnitary`. Their Fourier formulas are proved by `torusFourierEquiv_continuumTorusUnitary_apply` and `torusFourierEquiv_centralTorus`

`Unitary_apply`. The equality `norm_fullTorusResidual_eq_coeff` transfers the residual norm exactly. The final physical convergence statement is `fullTorusPropagator_compactTime_converges`, with the same mass, time-window, state, tolerance, and positive-spacing quantifiers as the coefficient theorem.

5 The unbounded coefficient Hamiltonian

Block 48 deliberately changes operator type. In namespace `KomorebiPhysics.Relativistic3DUnboundedDirac`, `diracCoeff` applies the continuum mode matrix at each coefficient. A square-summable input need not have a square-summable output after this operation. The maximal domain is therefore defined by the output condition

$$\mathcal{D}(H_m) = \{f \in \ell^2(\mathbb{Z}^3;\mathbb{C}^4) : (H_m(n)f_n)_n \in \ell^2\}.$$

The declaration `continuumDiracDomain` makes this set a complex submodule. The definition `continuumDiracOutput` packages the output only after domain membership has been supplied. The resulting partial linear map is `continuumDiracPMap`.

An input of subtype `(continuumDiracPMapm).domain` includes both its coefficient field and the proof that the output is square summable. This is the formal counterpart of writing $f \in \mathcal{D}(H_m)$ in an unbounded-operator theorem. Omitting that subtype would change the mathematical claim.

5.1 Density, symmetry, and the adjoint

The declarations `single_mem_continuumDiracDomain` and `coeffTruncation_mem_continuumDiracDomain` put all finite-support coefficient fields in the domain. Together with the earlier finite-tail approximation theorem, they establish `continuumDiracDomain_dense`.

The theorem `continuumDiracPMap_isFormalAdjoint` proves symmetry on the specified domain. The route to self-adjointness then has two separate inclusions. Symmetry and density give `continuumDiracPMap_le_adjoint`. To obtain the reverse inclusion, the proof tests an adjoint-domain vector against fields supported at a single Fourier label. The declaration `continuumDiracAdjoint_apply_coord` identifies each coefficient of the adjoint output with the same Dirac matrix action.

Since the adjoint output is itself an element of coefficient ℓ^2, this coordinate identity places the input in the maximal domain. That is `continuumDiracAdjoint_domain_le`. The theorem `continuumDiracAdjoint_le` then establishes the reverse partial-map inclusion. Equality with the adjoint is recorded in `continuumDiracPMap_selfAdjoint`, and closedness follows in `continuumDiracPMap_isClosed`.

The argument does not replace the adjoint domain by the original domain by definition. Its reverse inclusion is a proved statement. This is the essential distinction between a symmetric expression and a self-adjoint realization of an unbounded operator.

5.2 Finite Fourier approximation in the graph

The theorem `continuumDiracPMap_coeffTruncation` proves that truncation commutes with the multiplier. The statement `exists_finite_graph_approx` then takes $f \in \mathcal{D}(H_m)$ and $\varepsilon > 0$ and produces one finite set S satisfying both

$$\|f - P_S f\| < \varepsilon, \qquad \|H_m f - H_m P_S f\| < \varepsilon.$$

The same truncation controls the state and its operator image. Separate approximations with unrelated supports would not state this result.

This is the finite Fourier graph-approximation property underlying the graph-core interpretation. It does not assert graph convergence of a different family of lattice generators. Here the operator is fixed, and its domain vectors are approximated by finite Fourier fields.

6 The exact weighted domain and physical Hamiltonian

Block 49 works in namespace `KomorebiPhysics.Relativistic3DWeightedDomain`. Its first declarations expose the energy hidden in the maximal-domain definition. The function `modeEnergySq` is

$$E_m(n)^2 = m^2 + |k_n|^2.$$

The theorem `continuumHilbertModeOperator_sq` transports the earlier Clifford-square identity to the Hilbert mode operator. Combining that identity with modewise symmetry gives `norm_continuumHilbertModeOperator_sq`:

$$\|H_m(n)v\|^2 = (m^2 + |k_n|^2)\,\|v\|^2 .$$

The function `weightedModeTerm` is exactly the right-hand side with $v = f_n$. Thus `mem_continuumDiracDomain_iff_summable_weighted` identifies the maximal domain with summability of the modal energy.

6.1 Why the domain is independent of mass

The ambient type has already guaranteed

$$\sum_{n\in\mathbb{Z}^3} \|f_n\|^2 < \infty.$$

For any fixed real mass, the series $m^2 \sum_n \|f_n\|^2$ is therefore finite. The remaining issue is the momentum-weighted series. The source formalizes this through `summable_weighted_iff_momentum` and `summable_h1_iff_momentum`.

The definition `h1ModeTerm` uses the inhomogeneous weight $1 + |k_n|^2$. Consequently, `mem_continuumDiracDomain_iff_summable_h1` proves

$$f \in \mathcal{D}(H_m) \quad \Longleftrightarrow \quad \sum_{n\in\mathbb{Z}^3} (1 + |k_n|^2)\,\|f_n\|^2 < \infty.$$

This is the exact first-order Fourier condition. Calling it the first-order Fourier-Sobolev domain describes that explicit characterization. The formal statement uses `Summable` and `h1ModeTerm`, rather than an identification with every possible external implementation of a Sobolev space.

The theorem `mem_continuumDiracDomain_iff_mass` states equality of membership for any two real masses. It covers zero mass as well. The conclusion concerns the domain, not equality of the two Hamiltonians or their energy values.

6.2 Conjugation includes the domain

The physical domain is defined by `physicalDiracDomain` as the preimage of the coefficient domain under the Fourier equivalence. The map `physicalToCoeffDomain` retains the membership proof while

transporting a physical-domain vector into the coefficient domain. The definition `physicalDiracPMap` then gives the physical partial operator

$$H_m^{\mathrm{phys}} = \mathcal{F}^{-1} H_m \mathcal{F}, \qquad \mathcal{D}(H_m^{\mathrm{phys}}) = \mathcal{F}^{-1}\mathcal{D}(H_m).$$

This formula refers to a partial operator with a transported domain. It is not an everywhere-defined bounded conjugation.

The declarations `torusFourierEquiv_physicalDiracPMap_apply` and `physicalDiracPMap_apply_coeff` prove exact intertwining. The theorem `mem_physicalDiracDomain_iff_summable_h1` gives the physical statement

$$f \in \mathcal{D}(H_m^{\mathrm{phys}}) \quad \Longleftrightarrow \quad \sum_{n\in\mathbb{Z}^3} (1 + |2\pi n|^2) \left\|\hat{f}(n)\right\|^2 < \infty.$$

All Fourier coefficients here are those of the checked torus Fourier equivalence.

6.3 Self-adjointness in physical space

Density is `physicalDiracDomain_dense`. Symmetry is `physicalDiracPMap_isFormalAdjoint`. The proof of `physicalDiracAdjoint_domain_le` Fourier-transforms a vector in the physical adjoint domain and constructs the corresponding coefficient adjoint witness. The coefficient-domain theorem from Block 48 then applies. Fourier transport returns the resulting membership proof to physical space.

This argument produces the reverse adjoint-domain inclusion needed for `physicalDiracAdjoint_le`. Combining it with `physicalDiracPMap_le_adjoint` proves `physicalDiracPMap_selfAdjoint`. The theorem `physicalDiracPMap_isClosed` records closedness.

Finally, `norm_physicalDiracPMap_sq_eq_tsum_weighted` states the exact energy sum for a physical-domain vector:

$$\left\|H_m^{\mathrm{phys}} f\right\|^2 = \sum_{n\in\mathbb{Z}^3} (m^2 + |2\pi n|^2) \left\|\hat{f}(n)\right\|^2.$$

The related coefficient-space identity is `norm_continuumDiracPMap_sq_eq_tsum_weighted`. Both statements require a domain vector. Their series are justified by the domain condition, not by a formal rearrangement of a potentially divergent expression.

7 Verification records and the trust boundary

The v0.49.0 production receipt records Lean 4.34.0-rc2, compiler commit

```
6a10ac8c22beadecabdbb0919c2b50214762f91d
```

and Mathlib revision

```
aeb72f1e9981b6d35b426b16dca7759be3e7ba06.
```

The recorded platform is Linux x86_64. The supplied third-party notice states that the compiler, Mathlib, and their dependencies are external pinned software, not included in the Lean source package.

Recorded gate	v0.49.0 result
Production module import-closure build	Pass, 3,064 jobs.

Recorded gate	v0.49.0 result
Clean-stage module import-closure build	Pass, 3,064 jobs.
Exact final source-archive replay	Pass, 3,064 jobs.
Warning-as-error source compilation	Pass.
Block 49 API checks	40 declarations checked.
Block 49 theorem and lemma axiom audit	30 declarations checked.
New negative controls	Three correctly rejected.
Immediate parent Lean files	905 compared and byte-preserved.
Final Lean/config identity	919 files, as recorded by the handoff.
Recorded cumulative inventory	147 active mathematical modules, 3,920 named declarations, 2,968 theorem declarations, 207 negative controls.

The final-archive replay record identifies the source archive with SHA-256

`413b107ea21130ce45ff6ae261263d3a326e63047dfe1d9a389543f0928891aa`.

It records 5,452 internal checksum checks without mismatches, no unsafe paths, and no CRC errors. These archive checks address packaging identity and integrity. They do not replace the proof checks.

For the thirty new theorem or lemma declarations, the recorded axiom set is limited to `Classical.choice`, `Quot.sound`, and `propext`. The new-source scan reports zero occurrences counted as `sorry`, `admit`, `native_decide`, `unsafe`, or a custom axiom. The scope of that new-source scan should not be enlarged into an assertion that every dependency was rebuilt and audited from first principles. The pinned compiler and imported library environment remain part of the trust boundary.

The negative controls try three invalid claims: that two physical mass domains differ, that zero lies outside the physical domain, and that the constructed physical operator is not self-adjoint. Their rejection checks the intended statements and hypotheses at specific boundaries. It is not an exhaustive search for all possible specification mistakes.

The production receipt also retains failed attempts. Dependency builds exceeding an external command window were resumed and were not counted as passes. An early physical-adjoint draft failed during Fourier inverse simplification and subtype coercion. An API generation script failed with a `NameError` and produced no evidence. Only the repaired, completed gates support the reported release result.

7.1 Which files govern the current status

The versioned files `Komorebi_Physics_v0.49.0_RESULT.md`, `Komorebi_Physics_Block49_FINAL_STATUS.json`, and the v0.49.0 production, fresh, and final-archive replay receipts state the historical Block 49 frontier. The outer complete-handoff README describes that checkpoint and supplies its historical replay commands.

Several generic files inside the cumulative source archive remain historical. Its `README.md`, `START_HERE.md`, and `HANDOFF_MANIFEST.json` describe Block 45. Its generic `RESULT.md` describes Block 48. Its `lakefile.toml` retains the package version 0.48.0. Its generic `declaration_inventory.json` contains 142 modules and 3,718 declarations through Block 43. These files should not be used alone to infer that the later verified blocks are absent.

Conversely, their older statements should not silently be rewritten into new verification evidence. This editorial edition leaves the inherited source unchanged and records the discrepancy. The main

aggregate `KomorebiPhysics.lean` imports through the Block 47 torus evolution module and does not import the Block 48 and 49 modules. The official Block 49 module target is therefore the relevant entrypoint for the current endpoint. The receipts reviewed here do not establish a fresh successful default-root build. Nor does this companion assert that an older root-build failure has been reproduced in the current editorial session.

8 Replaying the recorded endpoint

The outer complete-handoff README supplies the following commands. They are reproduced without adding a new verification procedure. Run them from an extracted source root after attaching the exact pinned portable Lean and Mathlib environment.

```
lake build KomorebiPhysics.Relativistic3DWeightedDiracDomain
lake env lean -DwarningAsError=true KomorebiPhysics/Relativistic3DWeightedDiracDomain.lean
lake env lean -DwarningAsError=true CheckBlock49API.lean
lake env lean -DwarningAsError=true AuditBlock49.lean
```

The first command builds the module import closure. The second checks the source with warnings treated as errors. The API file checks the forty new named declarations, and the audit file prints the axiom dependencies of the thirty theorem and lemma declarations. The recorded release procedure also includes the isolated negative controls, parent preservation, and archive integrity checks. The four commands above are the documented local entrypoints, not a claim that those additional gates have been rerun by this editorial edition.

The source ZIP alone is not a complete offline runtime. An offline replay requires the pinned compiler and all corresponding dependency sources and caches to have been restored beforehand. The preceding generator replay recovered the seven-part offline runtime, checked its recorded combined digest, and used its exact dependency pins. The runtime itself remains an external dependency rather than part of the publication ZIP. Substituting an arbitrary later Lean or Mathlib version would be a compatibility migration and would require its own validation record.

9 The checked strong generator

The module `KomorebiPhysics/Relativistic3DGeneratorBridge.lean` imports the preserved physical Hamiltonian and works with its actual coefficient and torus types. It does not replace the infinite mode space by a finite test array. Its namespace is `KomorebiPhysics.Relativistic3DGeneratorBridge`. The source contains thirty-seven named theorems and two definitions, in addition to a local scalar instance. Compiler-generated auxiliary declarations are not counted as separate mathematical results.

9.1 The full-domain argument

The central statements have the following meaning. Let $U_m(t)$ be the previously constructed physical torus group and let H_m be the existing partial Hamiltonian. Then

$$\left.\frac{d}{dt}\right|_{t=0} U_m(t)f \text{ exists in } L^2 \quad \Longleftrightarrow \quad f \in \mathcal{D}(H_m), \qquad \left.\frac{d}{dt}\right|_{t=0} U_m(t)f = -iH_m f.$$

The domain on the right is the explicit first-order Fourier domain. No second Hamiltonian moment is assumed. The source also proves strong continuity for every L^2 field, not just fields in the operator domain.

The proof begins with the derivative of the existing bounded exponential on one four-spinor mode. A mean-value estimate bounds each difference quotient by the norm of that mode's Hamiltonian output. Square-summable domination then promotes the coordinate derivatives to a derivative in the full coefficient Hilbert space. This proves the forward domain inclusion. For the converse, a continuous coordinate evaluation maps an assumed Hilbert derivative to its uniquely determined modal derivative. Multiplication by i recovers a square-summable Hamiltonian output. Thus differentiability cannot silently enlarge the maximal domain.

The same fixed Fourier equivalence transports the argument to physical fields, including their domain witnesses. Applying the group law to a shifted orbit gives domain invariance. Uniqueness of derivatives then gives Hamiltonian commutation. In particular,

$$U_m(t)\mathcal{D}(H_m) \subseteq \mathcal{D}(H_m), \qquad H_m U_m(t) f = U_m(t) H_m f, \qquad \frac{d}{dt} U_m(t) f = -i H_m U_m(t) f.$$

Using the opposite time gives the reverse domain inclusion as an ordinary consequence of the checked group law. The derivative is continuous, and its norm is the constant $\|H_m f\|$. The corresponding Hamiltonian-output norm is preserved by evolution.

9.2 Declaration correspondence

Declaration in the generator namespace	Mathematical role
`coeff_tendsto_zero`, `coeff_tendsto`	Vector-valued square-summable domination.
`mode_orbit_hasDerivAt`, `mode_increment_bound`	Exact modal derivative and its output-dependent bound.
`continuum_orbit_continuous`	Strong continuity on all coefficient L^2.
`continuum_orbit_hasDerivAt`	Full-domain strong derivative at every time.
`mem_domain_of_hasDerivAt_zero`	Converse generator-domain inclusion.
`differentiableAt_zero_iff_mem_domain`	Exact coefficient generator domain.
`physical_differentiableAt_zero_iff_summable_h1`	Exact physical differentiability criterion in Fourier weights.
`physical_deriv_zero`	Generator value $-iH_m$ for the actual partial operator.
`physical_domain_invariant`, `physicalEvolvedDomain`	Evolution with an explicit domain witness.
`physical_hamiltonian_commutes`	Domain-correct Hamiltonian commutation.
`physical_schrodinger_hasDerivAt`	Strong physical Schrödinger equation.
`physical_velocity_continuous`, `physical_graph_energy_conserved`	Continuous velocity and conserved output norm.

The declaration inventory records the complete list, source locations, and source digest. The four cutoff commutation results remain useful finite identities, but they are not substituted for the Hilbert-space derivative proof.

9.3 Replay and trust boundary

The preceding generator checkpoint records a 3,066-job import-closure build and a 3,064-job baseline target replay. Those job totals include dependency traversal and cached products. They do not mean that the compiler or all of Mathlib was rebuilt from source. The compiler and Mathlib pins are the same as those stated above. The current centered-extension checks are recorded separately below.

The generator source is checked with warnings treated as errors. Separate gates check its named API, print every named declaration's axiom dependencies, exercise positive API examples, and reject altered interfaces with the wrong sign, a missing imaginary unit, or an omitted domain witness. A deliberately inserted proof placeholder is also rejected. These are targeted mutation controls, not an exhaustive specification search. The axiom output is restricted to `propext`, `Classical.choice`, and `Quot.sound`.

A separate-path copy of the exact source is checked without importing the compiled generator module. An imported-environment `leanchecker` run replays the new compiled declarations in the same Lean kernel. This is a check against environment manipulation, not an independent proof assistant, and it does not replay all imported mathematics from an empty environment. The earlier unsuccessful proof attempts are retained separately from the accepted results.

The current replay instructions are in `formal/continuum/README.md`. They use the supplied runner and exact pinned external runtime, and write to a new verification directory. The historical generator and baseline receipts remain unchanged. The old aggregate project root is not used to claim a fresh verification of every historical physics module.

10 The checked centered approximation

The centered approximation extends the same coefficient Hilbert space and physical Fourier equivalence used above. Four modules supply the new proofs: `Relativistic3DCenteredConvergence`, `Relativistic3DModalPerturbation`, `Relativistic3DCenteredRates`, and `Relativistic3DCenteredResolvent`, all in the `KomorebiPhysics` namespace. Their source contains seventy-four explicit theorem declarations, fifteen definitions, and eleven abbreviations. The compiled inventory has 162 declarations from these modules, including generated auxiliaries. Those counts describe the implementation, not a number of independent physical laws.

10.1 Bounded operators and the full graph domain

For every nonzero real spacing h, `centeredCoeffOperator` is the bounded operator whose modal action is the actual centered Dirac symbol. Its self-adjointness and bound

$$\|H_h\| \le \sqrt{m^2 + 3|h|^{-2}}$$

are proved. The exact Clifford identity for the difference from the continuum symbol controls the modal error by the first-order Fourier energy. Square-summable domination then gives $H_h f \to H f$ for every $f \in \mathcal{D}(H)$ as h approaches zero through either sign.

The bounded-operator constructor assigns an auxiliary zero operator at $h = 0$. Its identification with the centered modal multiplier requires $h \neq 0$, and the convergence theorem uses a punctured neighborhood. This convention does not replace the continuum Hamiltonian by a bounded operator. The separately defined propagator retains its removable continuum value at zero spacing.

10.2 Actual resolvents and their domains

For every $z \in \mathbb{C}$ with $\operatorname{Im} z \neq 0$ and every nonzero real spacing h, the source constructs the continuum and centered resolvents as bounded maps. It proves both inverse identities, proves that the continuum resolvent has range in the maximal Hamiltonian domain, and establishes the norm bound $|\operatorname{Im} z|^{-1}$. The construction applies in both complex half planes.

The resolvent identity is then an equality between the constructed maps, with the continuum domain witness supplied explicitly. Its norm estimate reduces the resolvent error to the graph-domain error already proved. Consequently the resolvents converge strongly on every coefficient vector, and Fourier conjugation gives the same conclusion on physical torus fields. No operator-norm resolvent convergence is asserted.

10.3 The third-order moment and the quadratic rate

The proposition `HasThirdOrderMoment` is the summability condition for the chapter's actual weight $\sum_j |k_{n,j}|^6$. It implies membership in the maximal Hamiltonian domain. The scalar sine remainder gives the coefficient $1/6$, while a modal interpolation argument bounds the propagator difference without assuming that the distinct centered and continuum Hamiltonians commute.

The checked global bounds are

$$\|(H_h - H)f\| \leq \frac{h^2}{6} M_3(f), \qquad \|(U_h(t) - U(t))f\| \leq \frac{|t|h^2}{6} M_3(f).$$

The generator estimate requires $h \neq 0$. The propagator statement covers every real time, including negative time. For $T \geq 0$, one envelope $Th^2 M_3(f)/6$ holds at every $t \in [-T, T]$. The physical-space results use the same moment of the actual Fourier transform. The source does not infer moment finiteness merely because Lean gives a total value to an infinite sum.

10.4 Declaration correspondence and accepted checks

The prefixes C, M, Q, and R below refer respectively to the centered-convergence, modal-perturbation, centered-rates, and centered-resolvent namespaces just named.

Declaration	Mathematical role
C: `centeredCoeffOperator_selfAdjoint`, `centeredCoeffOperator_norm_le`	Actual bounded centered operator and its global norm bound.
C: `centered_graph_tendsto`, `physical_centered_graph_tendsto`	Two-sided punctured limit on the complete maximal domain.
R: `continuumCoeffResolvent_mem_domain`, `continuum_shift_resolvent`, `continuum_resolvent_shift`	Constructed continuum inverse with the correct domain.
R: `centered_shift_resolvent`, `centered_resolvent_shift`	Both centered inverse identities for nonzero spacing.
R: `centered_resolvent_error_le`, `centered_resolvent_tendsto`	All-nonreal resolvent estimate and strong limit.
M: `modal_perturbation_bound`	Signed-time estimate without cross-commutation.
Q: `HasThirdOrderMoment.mem_continuumDiracDomain`, `centered_generator_quadratic_bound`	Actual weighted-domain witness and generator rate.

Declaration	Mathematical role
Q: `centered_propagator_quadratic_bound`, `centered_propagator_compactTime_bound`	Global and compact-time propagator rates.
R and Q: physical resolvent and rate theorems	Transport through the existing torus Fourier isometry.

The accepted run freshly compiled all twenty-six local modules in the required import closure. Its sixty-four recorded commands include a recursive axiom audit of 209 environment declarations: 162 from the four new modules and 47 from the unchanged generator module. The larger environment count includes generated auxiliaries. Only `propext`, `Classical.choice`, and `Quot.sound` occur as axiom dependencies.

Twenty-one controls have their expected outcomes: eight positive uses, twelve elaboration rejections, and one custom-axiom policy rejection. They exercise signed spacing and time, nonreal spectral parameters, the moment and domain hypotheses, and the distinction between an auxiliary zero-spacing definition and an actual inverse. Five module kernel replays cover the new modules and the generator.

The exact acceptance receipt is `formal/continuum/verification/block51/current/receipt.json`. It binds the source snapshot, commands, logs, compiled objects, declaration inventory, and relevant runtime identities. The compiler and external compiled dependencies are pinned and reused. The module replays use the same Lean kernel and imported environment, not a second proof assistant or a source rebuild of all dependencies. Failed development and verification attempts remain in the delivery record.

11 Remaining mathematical and physical comparisons

The chapter *The Dirac Group and Its Generator* proves the spectral-exponential identity by the spectral theorem and uniqueness of a strongly continuous unitary group with the specified generator. The checked full-domain derivative and generator value supply the central identification, but the abstract spectral-functional-calculus expression itself has not been linked to a separate Lean spectral-calculus API. No such formal linkage is claimed here.

The checked centered Fourier family acts on a fixed infinite coefficient space. A changing finite spatial lattice remains a separate comparison problem requiring sampling, interpolation, aliasing control, and compatible domains. The graph, resolvent, and third-order moment theorems do not supply those identifications by themselves.

Regular-connection link estimates also remain to be transported into a complete background-coupled Hamiltonian and evolution. Dynamical gauge fields require additional degrees of freedom and a Gauss constraint. Gravitational dynamics, cosmological production, and empirical predictions do not follow from the checked free-matter construction. Declaration totals do not measure completion of those physical problems.

12 Historical block index

The index below is a compact navigation aid for the cumulative source. It is based on the block proof guides and source maps. Where an individual guide is absent, it uses the corresponding source module description and inventory origin. An entry identifies the work undertaken at a block, including

corrections and explicit hypotheses. It is not a claim that every subject named in the historical master is fully formalized.

Block	Recorded contribution
1	Finite projector atlases, state and operator reconstruction, packet currents, coarse balance, and exact hypothesis-sensitive examples.
2	Moving frames and whole-object deformation under explicit differentiability and invertibility assumptions.
3	Nonlinear flow identities along supplied trajectories, refinement towers, mixed-state transfer, and source balance.
4	Finite tensor atlases, reduced states, partial traces, entropy and moment derivatives under their stated positivity hypotheses.
5	Lindblad jump formulas, coherence loss, population balance, and detailed-balance statements.
6	Exact refinement cancellation and analytic error-control estimates.
7	Positive, overlapping, and moving effect atlases.
8	Parseval channels, finite dilation constructions, and current gauge.
9	Explicit domains, countable continuity, and a genuine unbounded spectral model.
10	Bounded spectral operators and topology-specific truncation.
11	Constructed spectral solutions and the differentiability domain.
12	Infinite coarse cells, domain control, and current regrouping.
13	Finite characters, infinite residue masks, and maximal ladder domains.
14	Domain-correct resolvent commutators and nonreducing packets.
15	Infinite reconstruction in strong, compact, and Hilbert–Schmidt settings.
16	Finite pointwise derivation no-go and resolution through corner algebras.
17	Exact one-dimensional discrete relativistic sector.
18	One-dimensional shell asymptotics, exact group velocity, and a correction certificate.
19	Exact one-dimensional continuum dispersion and group-velocity limit.
20	Three-dimensional Clifford walk and doubled spectrum.
21	Three-dimensional continuum correction, velocity, and anisotropy.
22	Moving-momentum Lorentz-defect calculation.
23	Four-spinor Clifford transport with genuine unit-circle phases. The current inventory records this source as the U(1) repair.
24	Exact refinement towers and an abstract complete-metric limit from summable reconstruction defects.
25	Hilbert refinement with corrected unit-circle phases and a concrete exact cylinder tower.
26	Geometric subdivision of an edge into two half-edges and the associated directional compatibility.
27	Cubical refinement, local path compatibility, and curvature-controlled gluing defects.
28	Full three-axis geometric refinement of the ordered walk and a continuous-linear product-defect calculus.
29	Separation of straight-edge compatibility from curvature, including a flat-but-incompatible counterexample.
30	Continuous-linear finite-lattice realization and separate edge and curvature budgets.

Block	Recorded contribution
31	Euclidean Hilbert norm realization, constant-one directional mismatch estimate, and a sufficient dyadic summability regime.
32	Full seven-factor Hilbert defect bound and the independent cubical path-swap contribution.
33	Affine abelian connection links, exact affine plaquette formulas, and quadratic curvature bounds.
34	Coordinate-regular abelian connections, exact line-integral composition, and quadratic derivative-envelope bounds.
35	Unit-volume finite-cell interpolation into L^2, adjoint averaging, and best approximation.
36	Weighted octree reconstruction and density in the address-continuum L^2 space.
37	Cartesian binary coding of the octree continuum.
38	Binary pushforward measure and Euclidean cube volume.
39	Adjoint-transported Euclidean levels and strong convergence of orthogonal projections.
40	Concrete dyadic boxes, almost-everywhere cylinder identification, and equality of concrete and transported levels.
41	Fixed-mode centered Dirac generator consistency and exact symbol-square identities.
42	Finite Fourier packet orthogonality, Parseval identities, and L^2 generator consistency.
43	Bounded mode exponentials, group identities, and fixed-time finite-packet convergence.
44	Hilbert mode self-adjointness and exact propagator unitarity.
45	Compact-time quadratic bounds for fixed modes and finite packets.
46	Full coefficient-space unitary evolution and compact-time strong convergence for every square-summable field.
47	Surjective torus Fourier equivalence and physical-space evolution and convergence.
48	Maximal unbounded coefficient Hamiltonian, self-adjointness, closedness, and finite Fourier graph approximation.
49	Exact first-order weighted Fourier domain, mass independence, and the self-adjoint closed physical torus Hamiltonian.

The earlier unbounded spectral and resolvent constructions are useful mathematical infrastructure. Their presence does not automatically supply the later missing identification for the specific free Dirac operator. Reuse requires an explicit theorem connecting the particular types, domains, and hypotheses. Keeping that connection visible is the purpose of the source-to-mathematics crosswalk.

The Gauge and Refinement Formalization

Exact declarations, geometric hypotheses, and retained evidence

1 What this companion verifies

The formal gauge program proceeds from unit-circle transport to geometric refinement, separates edge reconstruction from curvature, and then supplies a concrete Hilbert space in which continuum reconstruction can be discussed. These are distinct achievements. A local gauge-covariance identity, an exact fine/coarse intertwining, a finite operator bound, and a strong limit of spatial projections do not assert the same thing.

This companion explains the selected source modules through Block40. Its central reading route is Blocks25–34; Blocks35–40 supply the reconstruction interface. The surrounding physics project has reached Block49, whose free torus Dirac result is treated in the Dirac monograph. The later free result does not automatically incorporate the background connection developed here.

The source is the retained checked Block49 archive. The accompanying record made a new comparison of source bytes, declaration locations and inherited evidence. It did not invoke Lean. All eighteen selected modules agree both with their archived members and with the inherited source digests used below. That source-identity result is stronger than recognizing a filename, and weaker than a new kernel check.

The evidence comparison also found an inherited inconsistency. One Block27 cumulative-build log contains a Git failure although its receipt records a successful empty log. Six reported build or compile logs from Blocks36 and40 could not be located in the retained archive. Those exceptions are preserved and explained in Section 12. They are not silently repaired or counted as newly verified passes.

1.1 How to use the source map

The full crosswalk is `editorial/gauge/lean_declarations.csv`. Each entry records its namespace, source file, line, source digest and block assignment. The corresponding JSON, `editorial/gauge/lean_source_map.json`, records the original ZIP member, inherited receipt and result of each identity comparison. The CSV contains 1,135 named inventory entries, including 705 theorem or lemma entries. The separate `lean_instances.json` records eleven additional named instance forms and three anonymous instance forms. These are source inventory counts, not numbers of independent physical laws.

The retained global `declaration_inventory.json` is an older 142-module snapshot. It is not used here to count the entire Block49 project. Its entries for the eighteen selected modules were checked individually against their actual source lines. Every selected entry matched. This distinction prevents a valid local index from becoming a misleading current global total.

Most short names below are relative to the namespace `KomorebiPhysics` followed by the module stem. A name such as `EdgeSubdivision.two_halfSteps_embed` additionally lies inside the indicated structure namespace. The CSV supplies the fully qualified name when the short name is ambiguous.

2 The reading map

The first three rows were integrated into the checked Block25 release. The corrected unit-circle transport originated as the Block23 topic; the complete-metric bridge was the Block24 topic. Their presence does not establish two additional independent release gates. The Block25 receipt separately records their theorem counts and the source identity of the integrated bridge.

Topic	Mathematical responsibility	Names	Results
23	Corrected unit-circle Clifford transport	67	38
24	Exact towers and summable metric defects	28	19
25	Hilbert preservation and the cylinder tower	84	51
26	Geometric half-edges and covariant interpolation	76	45
27	Cubical paths, face phases and Bianchi identity	89	52
28	Glued half-grid and finite-product defect calculus	82	56
29	Straight-edge compatibility and its counterexamples	65	42
30	Operator packaging and independent defect budgets	68	44
31	Genuine finite Hilbert norms and constant-one bound	59	44
32	Seven-factor scalar bound and two-stage comparison	40	29
33	Exact affine links and quadratic curvature estimates	85	48
34	Coordinate-regular background connections	92	58
35	Unit-volume cell reconstruction in genuine L^2	28	20
36	Weighted octree reconstruction and prefix density	47	23
37	Cartesian bit coding and quotient geometry	52	28
38	Binary measure and Euclidean cube volume	89	58
39	Transported Euclidean levels and projection limits	21	13
40	Concrete dyadic boxes and level identification	63	37
	Selected inventory total	1,135	705

Here “Results” counts both source theorems and lemmas, following the retained inventory convention. Named and anonymous instances are excluded from these two columns and recorded separately. In particular, the Block36,37 and38 release receipts use some different declaration-count conventions. The present table does not overwrite their original counts.

The first group of files has the stem `Relativistic3DCliffordTransportU1`, `LocalDiscreteGlobalContinuum` and `Relativistic3DCliffordHilbertRefinement`. Thereafter the stems follow the topics: `Relativistic3DGeometricEdgeRefinement`, `Relativistic3DCubicalRefinement`, `Relativistic3DFullWalkRefinement`, `Relativistic3DEdgeCompatibility`, `Relativistic3DCombinedDefect`, `Relativistic3DHilbertScalarDefect`, `Relativistic3DHilbertFullWalkDefect`, `Relativistic3DAffineConnection` and `Relativistic3DRegularConnection`. All live under `KomorebiPhysics/` with extension `.lean`.

3 Types fix the gauge convention

3.1 Unitary phases are actual unitary phases

The corrected local phase type is Mathlib's `Circle`. The earlier candidate used $\mathbb{C}^\times$, the group of nonzero complex numbers. Every element of $\mathbb{C}^\times$ is invertible, but multiplication by such an element need not preserve norm. Replacing that type by `Circle` is a mathematical repair, not a notational preference. The negative control `NonunitScalarIsometry.lean` records the sensitivity of the claimed isometry to this distinction.

The spinor is a four-component complex function. A site state is a function from `Site` to that spinor; a link and a gauge are functions from `Site` to `Circle`. The translation is an equivalence $\tau : \mathrm{Site} \simeq \mathrm{Site}$, so the inverse site appearing in the directional shift is supplied explicitly. Finite-site hypotheses enter the Hilbert sums, not the bare pointwise covariance identity.

The source link $u(x)$ is oriented from x to τx. Its gauge action is

$$u^g(x) = g(\tau x)u(x)g(x)^{-1}, \qquad (G_g\psi)(x) = g(x)\psi(x).$$

When the main monograph uses a backward transporter $U_{x,\tau x}$, the conversion is $u(x) = U^{-1}_{x,\tau x}$. One must convert the orientation together with the transformation law. Changing only the letter leaves the endpoint factors reversed.

3.2 Projector transport and covariance

A `CliffordInvolution` supplies a complex-linear involution A. Its channels are $P_\pm = (I \pm A)/2$. The module proves idempotence, complementarity and cross annihilation. Its directional transport is

$$(S_{A,\tau,u}\psi)(x) = u(\tau^{-1}x)P_+\psi(\tau^{-1}x) + u(x)^{-1}P_-\psi(\tau x).$$

The explicit inverse in `directionalShiftInv` is checked in both orders. The declaration `directionalShiftEquiv` packages the resulting reversible map. This pointwise equivalence is not yet a statement about an infinite-lattice operator domain.

The theorem `directionalShift_gauge_covariant` says exactly

$$S_{A,\tau,u^g}(G_g\psi) = G_g(S_{A,\tau,u}\psi).$$

After expanding the two channels, its proof ends with the following exact source fragment:

```
funext x
simp only [directionalShift, gaugeState, plus_phaseMul, minus_phaseMul,
  phaseMul_mul, phaseMul_add]
rw [gaugeLink_source_mul, gaugeLink_inv_target_mul]
```

The two cancellation lemmas encode the source and target endpoint factors separately. This is why the equality is a covariance theorem for the actual shift, rather than a desired transformation rule supplied as a field of a structure.

The same module treats local spin steps that commute with scalar phases and composes them with the spatial shifts. The full palindromic walk has the factor order $M - X - Y - Z - Y - X - M$. This occurrence order will later determine the coefficients in the scalar defect budget.

4 What the Hilbert and cylinder results add

The finite spinor pairing is conjugate-linear in the first input:

$$\mathrm{ip}(\psi,\phi)=\sum_{j=0}^{3}\overline{\psi_j}\phi_j.$$

The module `Relativistic3DCliffordHilbertRefinement` proves its sesquilinearity, positivity and definiteness, then proves that the actual Clifford channels are orthogonal when the involution is Hermitian. Summing over a finite site type gives `stateIP`. Reindexing through the site equivalence establishes preservation of this pairing by a directional shift.

This proof uses two inputs with different roles. Unit modulus controls the phases; Hermitian Clifford structure controls orthogonality of the two channels. Invertibility of the shift alone would not give the Hilbert conclusion. The spectral local mass step is treated separately, and composition gives the seven-factor preservation theorem.

4.1 Replication and its measure

The Block25 refinement changes Site to Site $\times$ Fin(2). The new child label is a spectator: translations and links are pulled back through the first coordinate, while the state is replicated. Exact walk intertwining follows for these lifted data.

Raw counting measure satisfies

$$\mathrm{stateIP}(I\psi,I\phi)=2\,\mathrm{stateIP}(\psi,\phi).$$

Thus replication is not an isometry for the unweighted fine sum. The half-cell measure removes exactly this factor. At depth n, the checked inner product is weighted by 2^{-n}. The negative control `RawReplicationIsometry.lean` prevents the raw and weighted statements from being confused.

The recursively lifted states and walks instantiate `ExactRefinementTower`. A reconstruction along the zero-child branch is exactly compatible with replication. In this realization the reconstructed sequence is constant, and its global state and evolution reduce to the base state and base walk. The theorem gives an exact cylinder model. No Euclidean continuum field is introduced by calling the tower a refinement.

4.2 The general metric bridge

The independent module `LocalDiscreteGlobalContinuum` defines `SummableDefect` for a sequence f_n in a metric space. Its fields give a summable real sequence b_n and the adjacent estimate

$$d(f_n,f_{n+1})\le b_n.$$

The theorem `SummableDefect.cauchySeq` derives the Cauchy property. In a complete target, `exists_limit` and `tendsto_limit` give an actual limit, and `dist_limit_le_tail` proves

$$d(f_n,f_\infty)\le\sum_{m=0}^{\infty}b_{n+m}.$$

The target space, the reconstruction and its adjacent estimate must be supplied by an application. The metric theorem does not identify that limit with a Dirac propagator or solve a gauge-field equation. Later connection budgets use precisely this conditional bridge.

5 Geometric half-edges replace spectator children

Block26 uses `HalfSite Site = Sum Site Site`. The left injection denotes a coarse endpoint and the right injection the midpoint on its outgoing edge. The half translation sends the endpoint to its midpoint and the midpoint to the next endpoint. Its square is the lifted coarse translation, as proved by `halfShift_sq`.

An `EdgeSubdivision u` supplies two links with the ordered contract

$$u_2(x)u_1(x) = u(x).$$

The interpolation is the actual transported midpoint value:

$$(I\psi)(\mathrm{inl}\, x) = \psi(x), \qquad (I\psi)(\mathrm{inr}\, x) = u_1(x)\psi(x).$$

Endpoint and midpoint gauges are independent. The first half-link transforms with the midpoint gauge on the left; the second transforms with the next coarse endpoint gauge on the left. Their internal midpoint factors cancel in the product. The theorem `EdgeSubdivision.embed_gauge_covariant` checks the interpolation under this action.

The declaration `EdgeSubdivision.two_halfSteps_embed` proves

$$S_{\text{fine}}^2 I = I S_{\text{coarse}}.$$

It uses the full Clifford channel split. It is not a scalar-translation identity relabelled as spin transport. The named `intertwiningDefect` is subsequently proved zero for these compatible data.

The half-cell pairing again has weight $1/2$. The theorem `embed_halfCellIP_isometry` proves exact normalization, and `embed_restrict_adjoint` identifies restriction as the Hilbert adjoint: average the endpoint with the midpoint transported back by the first sublink. The composite interpolation after restriction is idempotent and fixes interpolated states.

5.1 The one-axis massive walk has a precise factorization

A `LocalStepSubdivision` is additional data whose local half-step squares to the coarse local step. The theorem `refinedOneAxisWalk_embed` then compares $M - S - M$ with its factorwise subdivision

$$M_{1/2}^2 - S_{1/2}^2 - M_{1/2}^2.$$

This is not a theorem equating a coarse walk with the square of an arbitrary whole fine walk. Moving factors past one another would require additional commutation relations. The Hilbert-certified version retains the half-cell pairing throughout.

6 Cubical curvature and the exact obstruction

Block27 constructs one Boolean cube with eight vertices and twelve positively oriented edges. `CubeLinks` stores four edges for each axis. Its six face plaquettes are ratios of the two same-endpoint paths on the corresponding face. In the source convention,

$$\Omega_{xy}(z) = \frac{P_{xy}(z)}{P_{yx}(z)}.$$

The vertex gauge factors telescope along each path, so these Abelian plaquettes are gauge invariant. For each face, plaquette one is equivalent to equality of the two path phases.

The six origin-to-opposite-corner path phases are named `pathXYZ`, `pathXZY`, `pathYXZ`, `pathYZX`, `pathZXY` and `pathZYX`. Adjacent permutations differ by one face phase. Two transposition sequences yield the exact cube relation

$$\Omega_{yz}(1)\Omega_{xz}(0)\Omega_{xy}(1) = \Omega_{xy}(0)\Omega_{xz}(1)\Omega_{yz}(0).$$

This is `CubeLinks.cube_bianchi`. The factors commute because the formal link group here is $U(1)$. A non-Abelian Bianchi statement would require basepoint transports and ordered multiplication; the displayed scalar identity is not that theorem.

6.1 Flatness is local path independence

`CubeLinks.Flat` requires all six face phases to be one. The theorem `CubeLinks.Flat.all_corner_paths` then identifies all six corner transports. The derived theorem `CubeLinks.Flat.xy1_of_other_five` supplies the omitted upper xy face from the other five. These are single-cell assertions, not a global trivialization on an arbitrary lattice with nontrivial loops.

The chosen reference path to each child vertex defines a gauge-covariant interpolation. Since eight unitary copies of one spinor are used, `cubeIP` carries weight $1/8$. The theorem `CubeLinks.embed_cubeIP_isometry` checks this exact cell normalization. The finite-site version `cellEmbed_isometry` then sums the cell identities.

When two path phases satisfy $p = \Omega q$, the path-defect theorem identifies

$$p\psi - q\psi = q(\Omega - 1)\psi, \qquad \mathrm{ip}(D, D) = |\Omega - 1|^2\,\mathrm{ip}(\psi, \psi).$$

The namespace `CubeLinks` contains both `pathDefect_eq_curvature_transport` and `pathDefect_ip_eq_curvature_weight`. A fixed nonzero basis spinor makes zero defect equivalent to trivial phase. The checked `nonflatCube_defect_nonzero` supplies a concrete counterexample to universal path independence.

7 A full half-grid and two independent defects

Block28 glues child cubes into Site×CubeVertex. Each coordinate half shift toggles its bit and carries the coarse site on the second half step. The phase frame r defines an interpolation $I_r\psi(s, v) = r(s, v)\psi(s)$. Canonical fine links insert the coarse link on the carrying half-step and conjugate by the frame phases at its endpoints.

The exact two-half-step identities hold in all three directions. Their composition proves factorwise refinement of the complete seven-factor walk. The fine pairing uses $1/8$ of the counting sum. These

conclusions use the constructed canonical links. Their existence does not say that an arbitrary choice of fine links satisfies the same reconstruction contract.

7.1 Compatibility is the actual straight-edge hypothesis

Block29 isolates this issue in `StraightCompatibleX`, `StraightCompatibleY` and `StraightCompatibleZ`. Each condition equates the ordered two-half-link product with the framed intended coarse link. The contracts are gauge covariant. Under them, `arbitraryRefinedSevenFactorWalk_frameEmbed` proves exact full factorwise intertwining for arbitrary compatible fine links.

Curvature and reconstruction compare different data. Curvature compares alternative paths at a fixed scale; compatibility compares the fine ordered product with its intended coarse edge. Two exact source witnesses prevent either condition from replacing the other:

- `exact_refinement_can_preserve_nonflat_plaquette` retains a nontrivial coarse plaquette in its two-by-two fine Wilson loop while the walk refinement is exact.
- `quarterPair_fineDefect_nonzero`, together with `flatQuarterCube_flat`, gives flat local data whose chosen half-edge product is incompatible with the prescribed coarse link.

Consequently `no_universal_flatness_implies_refinement` formally rejects a curvature-only zero-defect rule. This obstruction is a positive result of the formalization: it identifies the missing hypothesis instead of hiding it inside a norm estimate.

7.2 Finite-product telescoping preserves factor order

The continuous-linear calculus in Block28 defines the defect $FI - IC$ for an interpolation $I : X \to Y$. The two-factor identity retains the transported local defects. Iterating it over the finite ordered list gives `norm_fullDefect_le_budget` when all relevant factors are contractive. No factor is silently commuted past another.

Block30 packages actual straight-edge maps and a separate cubical path-swap stage. Its operator norm initially comes from the ambient finite function-space norm. Finite dimensionality gives continuity, but does not identify that norm with the Hilbert counting norm. The source guide states this limitation explicitly. The next block supplies the missing Hilbert realization.

8 The constant-one Hilbert estimate

Block31 defines

$$\mathrm{HilbertState}(\mathrm{Site}) = \mathrm{EuclideanSpace}_{\mathbb{C}}(\mathrm{Site} \times \mathrm{Fin}(4)).$$

The conversion from the earlier function representation is algebraically invertible, and `stateToHilbert_norm_sq` identifies the squared norm with the real part of `stateIP`. The operator norms below are therefore actual finite ℓ^2 operator norms.

For two links v, u, the shift difference splits into orthogonal Clifford channels. The exact energy formula, reindexing and unit modulus imply

$$\|S_v - S_u\| \le \max_x |v(x) - u(x)|.$$

The declaration `directionalShiftDefectHilbertCLM_opNorm_le_sup` records this constant-one estimate. It is independent of the number of sites. A triangle inequality that treated both channels as arbitrary overlapping vectors would lose this structure.

8.1 The genuine normalized fine defect

The map `halfStateToHilbert` contains a factor $1/\sqrt{2}$, encoding the half-cell measure. For a `FineEdgePair` whose ordered product is v, `fineEdgeDefectHilbertCLM_norm_eq_coarse` identifies the normalized fine-defect norm with the corresponding coarse-link mismatch norm.

Define $E(x) = v(x)u(x)^{-1}$. Unit modulus gives

$$|v(x) - u(x)| = |E(x) - 1|.$$

Thus the principal theorem `fineEdgeDefectHilbertCLM_opNorm_le_ratioSup` says

$$\|D_{\text{fine}}\|_{\ell^2 \to \ell^2} \le \max_x |E(x) - 1|.$$

Its statement explicitly requires a finite nonempty site type, a `HermitianClifford`, a site equivalence, the half-edge pair and the intended coarse link. The scalar supremum is not a curvature invariant; it measures reconstruction error. The quarter-turn witness has ratio error two.

9 The complete scalar budget and its scope

Block32 conjugates directional shifts into the finite Hilbert space, proves their norm preservation and applies the finite-product theorem there. The two mass factors are identical on the compared sides and assumed contractive. If $\delta_x, \delta_y, \delta_z$ are the three link-ratio suprema, then

$$\|D_W\| \le 2\delta_x + 2\delta_y + \delta_z.$$

The coefficients count occurrences in $M - X - Y - Z - Y - X - M$. The result is `norm_hilbertSevenFactorDefect_le_edgeBudget`. Its `hilbertSevenFactorPairs` compare effective reconstructed and prescribed links on a common finite Hilbert space. The theorem should not be read as an already constructed arbitrary global fine-grid Hamiltonian comparison.

Uniform phase multiplication on the full finite Hilbert state gives `hilbertPhaseOp`. For $p = \Omega q$, `norm_hilbertPhaseDefectOp_eq_curvature` proves the exact operator identity in norm,

$$\|P_p - P_q\| = |\Omega - 1|.$$

The selected five-face set is

$$\mathcal{F}_5 = \{xy(0), xz(1), yz(0), yz(1), xz(0)\}.$$

It omits $xy(1)$, which the Bianchi identity constrains. The exact source order of the path pairs is XYZ/YXZ, YXZ/YZX, YZX/ZYX, XYZ/XZY, XZY/ZXY. The operators form a specified product comparison; the sum is not declared to be the defect of every possible physical gluing construction.

The combined theorem applies a checked two-stage composition identity to the walk comparison and this cubical phase comparison:

$$\|D_{\text{total}}\| \le 2\delta_x + 2\delta_y + \delta_z + \sum_{f \in \mathcal{F}_5} |\Omega_f - 1|.$$

The declaration is `norm_combinedHilbertDefect_le`. Its exact zero sector requires both exact reconstructed links and flat cubical data. These are sufficient conditions for this specified comparison, with no claim that one condition implies the other.

10 Connections generate the links and the scale law

10.1 Affine backgrounds

Block33 defines actual coordinate line-integral phases for an affine real Abelian connection. Additivity of the edge angle gives exact ordered half-edge composition. Thus the straight-link budget vanishes for the generated links rather than being assumed small.

The constant curvature components give exact plaquette formulas

$$\Omega_{ij}(p,h) = e^{ih^2 F_{ij}}.$$

The source proves $|e^{it} - 1| \le |t|$, then applies it to the selected faces. The result is

$$\|D_{\text{total}}(h)\| \le h^2(|F_{xy}| + 2|F_{xz}| + 2|F_{yz}|).$$

The theorem `norm_affineGeneratedCombinedHilbertDefect_le` applies to the actual finite-site link functions and the specified cubical phase stage. The uniform background $A_y = Bx$ supplies a nontrivial plaquette at $B = \pi, h = 1$. Vanishing straight-edge error therefore does not force flatness.

10.2 Coordinate regularity is explicit data

Block34's `RegularConnection3` stores A_x, A_y, A_z, the six cross-coordinate derivative fields, actual one-dimensional `HasDerivAt` identities, and the continuity along coordinate lines needed by the integral proofs. `CrossDerivativeBounds` supplies six nonnegative uniform bounds. This interface should be read literally; it is not an unnamed blanket assertion that all fields are smooth in every required sense.

For example,

$$\theta_x(p,h) = \int_0^h A_x(p + te_x)\,dt, \qquad u_x(p,h) = e^{i\theta_x(p,h)}.$$

Interval-integral additivity proves exact half-edge composition. The finite-site equality `regularHalfProductLinkXOn_eq_coarse` has the following exact proof body:

```
funext s
exact linkX_half_product C h (pos s)
```

The analogous declarations hold for y, z. No quadrature approximation is hidden in these equalities.

The fundamental theorem of calculus bounds coordinate variation. Integrating across the transverse edge gives a quadratic plaquette envelope. If B_{ji} bounds $|\partial_i A_j|$, the checked bound is

$$|\Omega_{xy}(p,h) - 1| \le (B_{yx} + B_{xy})|h|^2,$$

and similarly for the other faces. The derivative envelope is

$$\mathcal{A}_\partial = (B_{yx} + B_{xy}) + 2(B_{zx} + B_{xz}) + 2(B_{zy} + B_{yz}).$$

The theorem `norm_regularGeneratedCombinedHilbertDefect_le` proves

$$\|D_{\text{total}}(h)\| \le h^2 \mathcal{A}_\partial.$$

It assumes a finite nonempty site type, contractive mass maps, site equivalences, a position map, the regular connection and its derivative bounds, and a base point for the cube. The position map and site equivalences are not accompanied here by a theorem identifying an arbitrary global lattice with a continuum covariant derivative.

The plaquette estimate uses sums of absolute derivative bounds. It does not use a proved exact surface-integral formula for the curl of a general connection. Accordingly, it can be less sharp than a curl-only estimate. The non-affine example $A_y = \sin x$ has $F_{xy} = \cos x$, so this block extends beyond the constant-curvature affine class while retaining the stated sufficient envelope.

10.3 Summability is a derived scale statement

For $h_n = h_0 2^{-n}$, the actual bound has ratio $1/4$:

$$\varepsilon_n = h_0^2 \mathcal{A}_\partial \, 4^{-n}.$$

The declarations `summable_regularDyadicBudget` and the corresponding summability theorem for the generated operator norms establish the scale summability. `tendsto_regularConnectionLimit` and `dist_regularConnectionLimit_le_tail` then instantiate the complete-metric bridge when its reconstructed sequence satisfies the required adjacent bound. The rate follows from the quadratic connection estimate. A ratio $1/2$ would be a different bound, not the exact expression proved here.

11 From finite cells to a Euclidean reconstruction space

The final six modules establish a spatial interface. Their filenames end in `L2CellReconstruction`, `WeightedDyadicReconstruction`, `OctreeEuclideanCoding`, `BinaryLebesgue`, `EuclideanDyadicProjection` and `EuclideanDyadicBoxes`, each with the common `Relativistic3D` prefix.

11.1 Normalization precedes density

Block35's `UnitVolumePartition` is a finite measurable disjoint covering with every cell of measure one. Step interpolation is a linear isometry into genuine L^2. The averaging map is defined as its Hilbert adjoint and then identified with cell integrals. The composite is the orthogonal cell projection; Pythagoras and best approximation are proved for the actual maps.

Block36 replaces unit cell volume by a positive common volume v. Its normalized finite coefficient ψ_i corresponds to the continuum value $v^{-1/2}\psi_i$ on cell i. This factor is necessary for an isometry from unweighted finite coefficient space. A cell value and a normalized coefficient cannot be interchanged without changing the norm.

The octree address space is $\mathbb{N} \to \mathrm{Fin}(8)$ with uniform product probability measure. At prefix length n, there are 8^n cells of measure 8^{-n}. The theorem `Octree.finitePrefixSubmodule_dense` proves density

using the product measurable structure and L^2 induction. At this stage the continuum is the address space, not yet the Euclidean cube.

11.2 Coding is surjective, but not injective

Block37 splits one octant into three bits, proves that the finite and infinite product measures are preserved by this rearrangement, and decodes each bit stream by its binary series. The decoder is continuous and surjective onto the closed unit cube. Compactness and the Hausdorff property give a quotient map. The explicit two expansions of $1/2$ prove noninjectivity.

The resulting prefix-tail bound controls the spatial diameter of decoded cells. It does not by itself identify the pushforward measure with Lebesgue volume. Block38 supplies that additional argument: the one-dimensional distribution function is squeezed between adjacent dyadic fractions, yielding the uniform interval distribution; coordinate products and address rearrangement then give cube volume. The pullback U from cube L^2 to address L^2 is a linear isometry.

11.3 Transported levels are identified only in Block40

Block39 first defines closed Euclidean levels by taking the closure of the adjoint images of finite-prefix subspaces. Since $U^*U = I$, contractivity of U^* transports address approximants back to the cube. Nestedness and density imply strong convergence of the orthogonal projections for every fixed cube L^2 state.

Block40 identifies those levels with spans of concrete dyadic-box indicators. A prefix cylinder is contained in the decoder preimage of its closed box; equal measures upgrade that inclusion to equality almost everywhere. Closed boxes share boundary faces, so a pointwise disjointness assertion would be false. Equality of their L^2 indicator representatives is the relevant exact statement.

There is an index offset in this interface. The source's `initialLevelSubmoduled` uses the coordinate set $\{0, \dots, d\}$, hence prefix length $d + 1$. Its Euclidean level at index d corresponds to geometric depth $n = d + 1$. The monograph's volume formula 8^{-n} uses geometric depth. Applying it directly as 8^{-d} to the formal level index would be off by one refinement.

The resulting concrete projections converge strongly to the identity. This gives spatial approximation of each fixed state. It does not give an operator-norm limit uniformly over the full infinite-dimensional unit ball, nor does it prove convergence of the gauge-coupled generators or propagators. Those dynamical comparisons require compatibility and domain theorems beyond density alone.

12 Evidence, exceptions and reproduction

12.1 What was checked for the documented audit

The identity script `editorial/gauge/lean_collect_provenance.py` reads the extracted checked archive and edition root. It compares each selected source against its original ZIP member and the relevant inherited source digest, checks every selected declaration's source line, and resolves the log references in the selected receipts. It does not execute Lean or interpret a build log as a proof term.

All eighteen module-source comparisons passed, as did all 1,135 selected declaration-location checks. The script resolved 108 log references with matching inherited hashes. It also recorded one hash mismatch, six unresolved references and three available negative-control logs without an inherited hash.

These counts describe references in the selected receipts, including one explicitly excluded earlier attempt; they do not count 108 independent successful mathematical checks.

Blocks25–36 use block-specific reports under `verification/`. The selected Block33 report is `block33_production/report.json`; Block34 uses `block34_production_final2/report.json`. Blocks37–40 use their named `Production_Verification.json` reports. The generic `verification/final/report.json` is not treated as the gauge frontier. Full source and receipt identities are retained in the JSON map.

12.2 The inherited Block27 discrepancy

The Block27 report records a successful cumulative build with return code zero and the SHA-256 digest of an empty log. The retained file `verification/block27_final/cumulative_build.log` instead contains:

```
error: external command 'git' exited with code 128
```

Its digest therefore differs from the receipt. This log is excluded as evidence for that claimed successful cumulative gate. The report and the discrepant file are both preserved. The selected source digest, the direct Block27 module-compile log and the declaration-audit log do match their recorded identities. The discrepancy does not establish that the cubical theorem is false; it does establish that this cumulative gate cannot be authenticated from the retained log as presented.

12.3 Missing and unpinned records

Three Block36 references point to unavailable original paths: the cumulative-resume log, the import-closure build log and the direct warnings-as-errors compile log. Three Block40 build or compile references are also absent: `kp40_final_direct_compile.log`, `kp40_module_build2.log` and `kp40_cumulative_build2.log`. Their receipt claims remain inherited claims with incomplete attached evidence in this archive.

The three Block40 negative-control logs are present, but their receipt does not provide a comparison hash. The accompanying record records their current identities without retroactively treating those new hashes as original receipt authentication. Thus source identity is complete for the selected scope, while retained receipt completeness is mixed.

Earlier failed or interrupted runs are preserved in the original source archive. The Block37 report explicitly excludes an outer timeout before a terminal full-build result. Block38 excludes an earlier receipt with empty gates and a file-not-found error, as well as incomplete cumulative attempts. Block39 excludes a build-order failure and wrapper timeouts. These descriptions are not counted as successful checks. The later accepted receipts retain their own scope and trust boundary.

12.4 Reproduce identities first, then compile if needed

The edition contains the original ZIP under `recovered_sources/current/`. Extract it into a fresh working directory. From the edition root, the identity comparison can be repeated with:

```
python editorial/gauge/lean_collect_provenance.py EXTRACTED_LEAN_ROOT .
```

The first argument is the directory containing `lean-toolchain`, `lakefile.toml` and `KomorebiPhysics/`. The script rewrites the derived gauge provenance records and copies the retained receipts; it leaves the extracted Lean source untouched.

An actual compiler reproduction requires the pinned Lean `v4.34.0-rc2` environment and Mathlib revision

aeb72f1e9981b6d35b426b16dca7759be3e7ba06.

The archive is not a promise that those runtime binaries and dependency objects are already installed. Read its `AGENTS.md`, `CONTINUE_HERE.md`, block-specific proof guides and verifier before running a gate. Do not substitute a currently installed unpinned compiler and present the output as the original reproduction.

For example, after restoring the pinned environment, a targeted regular-connection build uses the source module itself:

```
lake build KomorebiPhysics.Relativistic3DRegularConnection
lake env lean -DwarningAsError=true \
  KomorebiPhysics/Relativistic3DRegularConnection.lean
lake env lean AuditBlock34.lean
```

These are reproduction instructions, not commands executed for the documented audit. A complete gate must also run its API checks and isolated negative controls, validate their intended diagnostics, and write a terminal receipt tied to the source digests. A negative control's nonzero exit is useful only when it rejects the intended false mathematical assertion; an unknown import or missing executable is not that result.

The inherited runs used the pinned compiler and Mathlib cache. They did not rebuild Lean and Mathlib from compiler source. Source-identity preservation of predecessor modules is an explicit part of the modular evidence chain. A claim of a wholly fresh from-source verification would exceed both those receipts and the present editorial work.

13 The next formal junction

The selected program now has exact unit-circle covariance, geometric edge and cubical constructions, independent reconstruction and curvature budgets, finite Hilbert estimates from regular background links, and a concrete Euclidean reconstruction space. The unresolved junction is their dynamical integration.

The next background-gauge result must specify the continuum operator, its domain, the finite-resolution embeddings, and the precise sense in which the discrete action approaches that operator. Density of cell states supplies approximants; the quadratic link budget supplies local control. Neither supplies the missing operator-domain identification by itself.

The historical continuation file lists background gauge coupling as a remaining target. The free theory now has checked strong-generator, centered graph and resolvent convergence, and third-order moment results. The abstract spectral-functional-calculus identification remains a written deduction without separate formal API linkage. Establishing dynamics for gauge-field degrees of freedom and a Gauss constraint is a further step beyond a prescribed background connection. The non-Abelian derivations in the monograph remain ordinary mathematical developments unless and until their own ordered transport statements receive a corresponding formal source and receipt.

The Gravity Verification Companion

Proof status, computation, and the remaining formal bridge

1 The actual verification frontier

The gravity program now has an explicit chain of finite and resolved constructions: linearized constraint propagation, nonlinear detail return, spectral constraint closure, a real Lorentzian parent, a local canonical reduction, and spherical strong-field models. Their hypotheses are part of the results. They do not merge into one unrestricted theorem merely because their formulas can be assembled in one volume.

The formal source frontier remains Block 49. Its newest result concerns the maximal free Dirac operator on the physical torus. The accompanying preparation record did not run Lean and did not add a gravity module. The gravity monograph presents ordinary mathematical derivations and selected computational evidence. The title of this companion reflects that distinction.

There is nevertheless a genuine formal interface. The checked source contains the pointwise-derivation obstruction, matrix corner derivations, curvature commutators, exact projector coarsening, and a complete-metric reconstruction theorem. Eight modules and twenty-one selected declaration locations are mapped below. These provide reusable mathematics. They do not yet instantiate an Einstein–matter constraint algebra, a Grassmann-valued BV charge, or a nonlinear gravity evolution tower.

A new parity audit also found a substantive error in the inherited Lorentzian BRST presentation. The archived dual-spinor formula puts an odd Lorentz ghost on the right with the commuting-parameter sign. For odd fermionic matter and a left odd differential, that sign fails nilpotency. Section 7 gives the exact counterexample and the corrected right-action sign. The old numerical covariance checks use commuting spinors and parameters, so they could not detect this error. Their retained pass labels are preserved as historical evidence rather than promoted into authentication of the printed master equation. A second review gate leaves the physical massive-Dirac convention matching open, as explained in Section 8.

1.1 Four meanings of a successful check

An ordinary proof establishes a mathematical implication under its stated assumptions. A source-identity check establishes that a cited file is the exact retained archive member. A bounded computational check establishes the assertion actually implemented for its inputs and arithmetic. A Lean kernel check establishes a formal declaration in a specified environment. The accompanying record supplies the first three kinds in selected scopes and cites inherited formal declarations. It supplies no new kernel receipt.

The durable crosswalk is `editorial/gravity/verification_source_map.json`. It records seven central claim groups, sixty-nine distinct archive members, eight formal modules and twenty-one declaration locations. Every member has a SHA-256 digest and its original ZIP path. The source map is a route back to the evidence, not a count of independent physical laws.

2 The central claim ledger

The technical supplement identifies C08–C14 as the nonlinear and strong-field chain. Its chapter numbers follow the compiled Volume I order. This companion uses the exact source filenames as the stable locators, because those filenames retain earlier numbering.

ID	Result family	Scope needed to read the result
C08	Homological BFV transfer	Nilpotent fine differential, admissible contraction, invertible return operator and a compatible Hamiltonian or cyclic structure.
C09	Resolved Ashtekar algebra	Band-limited fields and smearings, sufficient headroom $16K < N$, detail bands retained before aliasing.
C10	Lorentzian classical BV parent	Declared Einstein–Cartan–Maxwell–Dirac action and consistent graded gauge conventions. The dual-spinor sign needs correction; physical massive-Dirac convention matching remains open.
C11	Boundary BFV and time slab	Regular nondegenerate boundary reduction. ADM ghost coordinates require sufficient headroom $32K < N$. The time-slab relation has separate gauge-fixing and solution hypotheses.
C12	Second-class reduction and time step	Oriented regular time gauge, positive triad determinant, an invertible endpoint Dirac matrix, and an exactly solved implicit step.
C13	Charged spherical mass coordinates	Spherical electrovac on a monotone areal-radius patch. Strong-field reconstruction has additional square-root and derivative lower bounds.
C14	Scalar flux and trapped sphere	Spherical massless scalar in the declared flat-slice gauge. The formation benchmark uses regular excision and bounded numerical controls.

The inherited seven ledgers contain 121 rows labelled pass: respectively 9, 9, 30, 16, 26, 19 and 12. They also retain one conditional-theorem row and two open rows in the parent and boundary ledgers. These counts are a source inventory. Some rows explicitly record analytic arguments with a zero residual. Those rows were not numerical evaluations of a complete BV charge. A zero in such a ledger must be read together with its method and note fields.

The current edition adds a bounded reproduction of the underlying second-class verifier, nine strong-field identity and archive-consistency checks, and seven exact BRST-parity checks. These three scopes are distinct. The parity checks include a successfully detected failure of the archived sign, so “checks pass” does not mean every inherited mathematical assertion survived unchanged.

3 How the nonlinear route changed

The early linearized ADM construction uses commuting derivative operators and a quadratic action. On that model, divergence identities propagate the linearized Hamiltonian and momentum constraints. The argument is valuable precisely because its cancellation can be inspected directly. It does not prove that the same identities survive arbitrary nonlinear pointwise multiplication.

The nonlinear chapter exposes the obstruction. In a finite pointwise algebra, an exact scalar-linear Leibniz derivation must vanish. Ordinary finite differences avoid being zero by failing that Leibniz identity. Contracting the defect into a particular constraint bracket may produce a visible nonzero witness, but one such witness is narrower than the universal algebraic obstruction. The current monograph preserves both levels of argument.

The next construction changes the algebra. Matrix corner derivations obey an exact product rule, while homological transfer retains coarse-to-detail-to-coarse returns. The transfer theorem assumes a consistent fine nilpotent theory and contraction data. It does not construct a microscopic four-dimensional gravity theory from those assumptions.

The resolved Ashtekar construction supplies a more concrete classical input. Trigonometric-polynomial products are allowed to occupy larger support bands. Below Nyquist, multiplication, differentiation, trace and canonical variation agree with the continuum polynomial calculation. This produces exact resolved bracket identities without making the complete finite pointwise algebra into a differential algebra.

The later Lorentzian chapters use a real first-order parent rather than imposing physical reality only after choosing one complex self-dual chart. The boundary construction then splits time and space. Its ledger still records second-class reduction as open at that stage. The following chapter provides a local regular time-gauge reduction. The earlier open row is therefore historical, while global reduction and degenerate strata remain open.

The discrete-gradient theorem supplies a finite-dimensional invariant-preserving step once a regular gauge patch and an exact step solution are available. The spherical electrovac and scalar chapters give additional reduced models and benchmarks. None of these later developments retrospectively proves a global Einstein–matter continuum limit or a quantum anomaly-cancellation theorem.

4 What is already formal: the algebraic obstruction

The namespace `KomorebiPhysics.PointwiseNoGo` defines the pointwise algebra $\iota \to \mathbb{C}$ and a complex-linear map D with the actual Leibniz property. Its theorem `finite_source_no_go` assumes a finite decidable index type and concludes $D = 0$. The stronger source theorem `derivation_eq_zero` uses coordinate idempotents without needing a finite-type hypothesis. Neither theorem concerns every possible finite-dimensional commutative algebra.

The proof is short enough to explain. A coordinate idempotent e_a satisfies $e_a^2 = e_a$. Leibniz gives

$$De_a = 2e_a De_a.$$

At coordinate a, this says $De_a(a) = 2De_a(a)$; elsewhere it says $De_a(b) = 0$. Thus $De_a = 0$. Applying the rule to $e_a f = f(a)e_a$ forces $Df(a) = 0$ for every a. This is the finite pointwise obstruction used by the gravity discussion.

The namespace `KomorebiPhysics.CornerDerivation` changes the product space to finite complex matrices and defines

$$\delta_K X = i(KX - XK).$$

The theorem `delta_mul` proves

$$\delta_K(XY) = (\delta_K X)Y + X(\delta_K Y).$$

For Hermitian K, `delta_preserves_hermitian` preserves Hermitian inputs. The theorem `trace_integration_by_parts` proves

$$\operatorname{tr}((\delta_K X)Y) = -\operatorname{tr}(X\delta_K Y).$$

These are actual algebraic identities for the source operators, rather than desired properties inserted into a structure.

In `KomorebiPhysics.CornerCurvature`, the curvature matrix is

$$F_{KL} = i[K, L].$$

The declaration `derivation_commutator` proves

$$\delta_K \delta_L X - \delta_L \delta_K X = \delta_{F_{KL}} X.$$

The commuting-generator hypothesis in `derivations_commute_of_generators_commute` is explicit. To use commuting corner derivatives in the linearized ADM proof, one must supply that hypothesis for the chosen generators. The existence of the theorem does not identify arbitrary matrices with physical spatial derivatives.

5 What is already formal: balance and refinement

The formal refinement results act on projector atlases, finite sums and currents. In `KomorebiPhysics.Atlas`, `coarsen_energy` adds packet norm squares over a fiber, and `coarsen_current` adds the corresponding fine current entries. These results come from `KomorebiPhysics/Refinement.lean`. They ensure that exact coarse aggregation preserves the established balance law.

The namespace `KomorebiPhysics.RefinementTower` proves composition of aggregation. Its declarations `push_comp`, `incidence_comp` and `aggregateCurrent_comp` show that direct and iterated sums or incidence-matrix transfers agree. The current transfer has the form $F \mapsto RFR^T$. This is exact path independence of that aggregation. It is not automatically path independence of nonlinear canonical gravity reductions.

The namespace `KomorebiPhysics.UniversalBalance` proves cancellation of internal antisymmetric fluxes. The theorem `source_boundary_sum` separates sources and boundary current; `hasDerivAt_total` removes internal transfer from the total derivative. To use this for spherical gravity, a separate derivation must identify the actual mass density, flux, differentiability and boundary conditions. The exact scalar mass-flux theorem remains that additional ordinary proof.

The strongest abstract continuum bridge here is in `KomorebiPhysics.LocalDiscreteGlobalContinuum`. An exact tower includes embeddings I_n, level evolutions U_n, and the assumption

$$U_{n+1} I_n = I_n U_n.$$

The declaration `ExactRefinementTower.evolve_iterEmbed` propagates this relation through successive levels. A complete metric reconstruction then uses adjacent bounds

$$d(f_n, f_{n+1}) \le b_n, \qquad \sum_{n=0}^{\infty} b_n < \infty.$$

The declarations `SummableDefect.exists_limit` and `dist_limit_le_tail` give convergence and the tail estimate

$$d(f_n, f_\infty) \le \sum_{j=0}^{\infty} b_{n+j}.$$

The theorem `ExactRefinementTower.ContinuumRealization.globalEvolution_error_bound` applies the same mechanism after levelwise evolution.

Gravity still needs its own state spaces, reconstruction maps, regular chart control, and summable evolution defects. A few-grid convergence fit does not supply that summable tower, and the abstract theorem does not construct it. This is a precise formal gap rather than an absence of useful formal progress.

6 The spectral headroom gate

For the resolved vacuum algebra, a field in $\mathcal{T}_K$ has coordinatewise Fourier support at most K. Products satisfy

$$\mathcal{T}_K \mathcal{T}_L \subseteq \mathcal{T}_{K+L}.$$

If a product has support below $N/2$, its sampled derivative agrees with the derivative of the trigonometric polynomial. The normalized grid trace is its exact zero Fourier mode. The rule is checked for the product actually used; a low support bound on each factor alone is insufficient.

The inherited C09 theorem counts the smearing among the band-K factors. Its largest canonical gradients have four factors, and a bracket has at most eight. The sufficient condition is therefore

$$8K < N/2, \qquad \text{or equivalently } 16K < N.$$

On this domain, the Hamiltonian bracket closes into the vector constraint, which is a diffeomorphism constraint plus a Gauss term. The densitized lapse has weight -1; omitting that weight changes the mixed bracket.

The retained C09 computation uses an 18^3 grid with band 1 inputs, so its declared sufficient headroom holds. It records seven bracket residuals of order 10^{-17}, one support check and one aliasing negative control. The band 3 control violates the hypothesis and gives an absolute Hamiltonian-bracket defect 3.288599×10^{-8}. This separates an observed resolved case from an observed aliased case.

A nonlinear Hamiltonian vector field can create new support. A fixed $\mathcal{T}_K$ is not generally invariant under evolution. The construction therefore uses a filtered tower and refines before products leave the allowed headroom. Exact bracket evaluation below Nyquist is distinct from existence of a global, invariant finite-band flow.

The parent and boundary audit grids need additional care. Their retained ordinary covariance checks use $N = 12, K = 1$ and a band 1 10^4 grid respectively. Those sizes do not satisfy the sufficient whole-parent $16K < N$ or whole-ADM-ghost $32K < N$ bounds. Individual lower-degree identities can have smaller support requirements. Those computations cannot be cited as complete numerical verification of the full master charge under the stated sufficient bounds.

7 The Lorentzian BRST parity correction

Fix a left odd derivation s, so for homogeneous quantities

$$s(ab) = (sa)b + (-1)^{|a|}a(sb).$$

Consider only the internal Lorentz spin sector. Let the ghost matrix χ and the row spinor $\bar{\psi}$ both be odd, with

$$s\chi = \chi^2.$$

The archived formula places the ghost on the right and uses

$$s\bar{\psi} = -\bar{\psi}\chi.$$

Its second derivative is

$$s^2\bar{\psi} = -((s\bar{\psi})\chi - \bar{\psi}(s\chi)) = 2\bar{\psi}\chi^2.$$

That expression need not vanish. The sign appropriate for a commuting infinitesimal Lorentz parameter cannot be copied unchanged after moving an odd ghost through odd matter.

An exact finite witness uses Lorentz spin generators $T_1 = \gamma_{23}/2$, $T_2 = \gamma_{31}/2$ for the mostly-plus Clifford convention. Their commutator is $-\gamma_{12}/2 \neq 0$. For independent odd generators θ_1, θ_2, η, set

$$\chi = \theta_1 T_1 + \theta_2 T_2, \qquad \bar{\psi} = \eta r, \qquad r\gamma_{12} \neq 0.$$

Then $\chi^2 = \theta_1\theta_2[T_1, T_2]$, and the archived sign gives a nonzero degree-three coefficient.

The consistent right-row convention is

$$s\bar{\psi} = +\bar{\psi}\chi.$$

It gives $s^2\bar{\psi} = \bar{\psi}\chi^2 - \bar{\psi}\chi^2 = 0$. Equivalently, one may keep the negative dual-representation sign with all odd ghosts placed on the left in component formulas. Under this convention the mass bilinear also obeys

$$s(\bar{\psi}\psi) = (s\bar{\psi})\psi - \bar{\psi}(s\psi) = 0.$$

The new exact check uses $\mathbb{Q}(i)$, seven Grassmann generators and actual 4×4 Clifford matrices. It verifies all 16,384 exterior-basis products and records seven passing checks, including the expected nonzero archived residual and the zero corrected residual. Its schema 2 manifest records a measured final run, actual input hashes before and after, and declared result hashes. This is a decisive finite counterexample to the printed pure-Lorentz nilpotency claim. It is not an independent verification of the complete diffeomorphism, electric, BV or boundary charge.

8 Parent and boundary claims after the correction

The real parent brings the tetrad, Lorentz connection, Maxwell auxiliary field and Dirac matter into one declared first-order action. Its ordinary commuting-parameter symmetry calculations remain useful evidence about curvature, torsion, covariance and the action density. The parity finding limits what follows from those checks: the printed graded differential needs correction before its nilpotency can support a classical master action.

There is a separate physical convention gate. The inherited signature is $\eta = (-,+,+,+)$ and the Clifford relation is $\{\gamma^I, \gamma^J\} = 2\eta^{IJ}$. The flat operators $i\gamma^I\partial_I \pm m$ have product $-\Box - m^2$, so either real-mass sign gives $(\Box + m^2)\psi = 0$. Plane waves therefore obey $\omega^2 = |k|^2 - m^2$, rather than the conventional positive massive shell. The covariance checks do not test that dispersion. The displayed gauge algebra can be studied as a declared model, but matching its factors of i, adjoint and mass sign to a physical massive Dirac action remains open. The free torus result from Block 49 does not repair this separate parent convention.

The classical BV construction has a clear dependency order. First establish a graded differential with $s^2 = 0$. Then establish $sS_0 = 0$ using the same parity and left/right conventions. Then insert the corresponding antifield terms and verify the chosen antibracket signs. An off-shell closed gauge algebra permits the usual minimal construction, but the actual component formulas must realize it consistently. The current monograph makes the dual-spinor repair explicit. The finite parity check certifies the local repair; it does not silently replace a full component audit.

The boundary construction has further assumptions. Its $3 + 1$ split identifies multiplier constraints and candidate moment maps. The full normal moment-map statement remains conditional on embedding data, a projectable regular quotient and equivariance. Obtaining a canonical boundary bracket requires a regular nondegenerate reduction. ADM ghost coordinates are a field-dependent change of ghost variables, so the canonical pullback must include the conjugate momenta and their induced terms. Sufficient support headroom belongs to this complete coordinate construction, not merely to one multiplication test.

The adaptive time-slab relation is a relation defined by the resolved variational problem. Its charge-intertwining statement depends on regular gauge fixing and an exact or suitably controlled solution of that problem. It is not automatically an explicit, single-valued update map. The following local discrete-gradient construction is a separate finite-step result.

The later second-class chapter pulls the charge back only if the BFV vector field is tangent to the regular second-class surface. That tangency condition is an actual theorem hypothesis. It is not removed by presenting a solution map for the algebraic constraints.

Finally, the quantum master equation has an additional measure or BV-Laplacian term. A cyclic classical transfer does not automatically supply an anomaly-free regulated gravity–matter measure. The inherited C10 ledger correctly leaves the quantum master equation open. The present parity correction concerns the classical starting point and gives no quantum anomaly-cancellation result.

9 The regular canonical reduction

The local reduction works in an oriented time-gauge sector with

$$E^0 = 0, \qquad e = \det(e_a{}^i) > 0.$$

The gravitational simplicity pullback removes the spatial rotational-connection term from the symplectic potential and leaves the boost connection $K_a{}^i$ paired with the densitized triad. The full antisymmetric internal-index sum is essential to the factor $1/\kappa$.

Maxwell auxiliary elimination is algebraic on the nondegenerate sector. The electric momentum pairs with the spatial potential. The fermionic primary matrix is invertible for $e \neq 0$. Defining the half-density spinor $\Psi = e^{1/2}\psi$ cancels the determinant-variation terms and gives the local reduced one-form

$$\Theta_{\text{red}} = \int_\Sigma \Big[\kappa^{-1}\widetilde{E}^a{}_i\,\delta K_a{}^i + \mathcal{E}^a\delta A_a + \frac{i}{2}(\Psi^\dagger\delta\Psi - \delta\Psi^\dagger\Psi)\Big]d^3x.$$

The inherited convention gives

$$\{K_a{}^i(x), \widetilde{E}^b{}_j(y)\}_D = \kappa\delta_a{}^b\delta^i{}_j\delta^3(x-y), \qquad \{\Psi_A(x), \Psi^\dagger_B(y)\}_D = -i\delta_{AB}\delta^3(x-y).$$

The torsion equation fixes the rotational connection algebraically and produces the axial-current interaction. In the parent signature and orientation conventions, the source gives a linear contorsion contribution $-3\kappa J_5^2/8$ and the total on-shell action correction $-3\kappa J_5^2/16$. Changing signature, orientation, gamma conventions or left/right fermionic derivatives requires rechecking those signs together.

The reduced chart still contains first-class rotation, electric, diffeomorphism and Hamiltonian constraints. “Minimal” means that redundant algebraic and fermionic-primary variables have been eliminated. It does not mean that all gauge orbits have been globally quotiented.

Degenerate triads are outside this regular chart. The inherited negative control sets $e = 0$ and obtains a rank-zero fermionic primary matrix. The local result therefore identifies a real boundary of applicability. It does not provide a stratified extension across that boundary or a global time-gauge section.

10 Exact time steps and approximate solves

At a finite resolved level, let $J^T = -J$ be the constant canonical matrix, let H be an autonomous Hamiltonian, and let Φ collect the constraints and regular gauge conditions. For endpoints z_0, z_1, the average-vector-field discrete gradient satisfies

$$F(z_1) - F(z_0) = \overline{\nabla} F(z_0, z_1)^T (z_1 - z_0).$$

Let G_Δ have the discrete gradients of the components of Φ as its rows. Assume the endpoint matrix

$$\Delta_\Delta = G_\Delta J G_\Delta^T$$

is invertible. Define

$$J_D^\Delta = J - J G_\Delta^T \Delta_\Delta^{-1} G_\Delta J.$$

Because Δ_Δ and its inverse are skew, J_D^Δ is skew. Direct multiplication gives $G_\Delta J_D^\Delta = 0$. Thus every exact solution of

$$z_1 - z_0 = \Delta t \, J_D^\Delta \overline{\nabla} H$$

preserves Φ and H. This implication is finite-dimensional linear algebra plus the exact discrete-gradient identity.

The approximate-solve distinction can also be quantified. If the implemented endpoint has residual

$$r = z_1 - z_0 - \Delta t \, J_D^\Delta \overline{\nabla} H,$$

while the discrete gradients and inverse are exact, then

$$\Phi(z_1) - \Phi(z_0) = G_\Delta r, \qquad H(z_1) - H(z_0) = \overline{\nabla} H^T r.$$

Consequently,

$$\|\Delta\Phi\| \le \|G_\Delta\| \|r\|, \qquad |\Delta H| \le \|\overline{\nabla} H\| \|r\|.$$

Quadrature and floating-point errors add their own terms. Small solver residuals are evidence for small invariant drift in a controlled patch; they are not exact preservation in finite arithmetic.

Local step existence uses a regular neighborhood and sufficiently small Δt. Long-time bounds require a coercive Hamiltonian and continued membership in the patch. An infinite tower additionally needs level-independent inverse-matrix, derivative and coercivity bounds. Those uniform gravity–matter estimates remain open.

11 The new bounded second-class reproduction

The accompanying record reran the underlying source program `Resolved_Second_Class_Reduction_Verification.py` in an isolated evidence directory, with unchanged source bytes. The seed is 20260811. The program combines symbolic torsion identities with floating-point pullbacks, Clifford identities, rank checks and one constrained-particle trajectory.

The trajectory is a six-dimensional surrogate with

$$q \cdot q = 1, \qquad q \cdot p = 0.$$

It uses 1,000 steps of size 0.1. It is not a simulated gravity–matter trajectory. The final run returned 26 passes, 4 information entries and 0 failures. The retained and new values agree. Representative values are

Bounded assertion	New value
Energy drift	7.771561×10^{-16}
Constraint drift	3.885781×10^{-15}
Maximum implicit-solve residual	1.092187×10^{-16}
RK4 energy-drift negative control	1.396105×10^{-6}
RK4 constraint-drift negative control	2.594613×10^{-6}
Gravity symplectic-pullback residual	1.110223×10^{-16}
Half-density identity residual	3.662053×10^{-15}
Exact on-shell contorsion coefficient	$-3/16$

The implementation thresholds are 10^{-10} for selected residuals and 10^{-8} for the negative-control drift. The exact symbolic coefficient is compared as an exact string after the symbolic calculation. The rank tests and determinant tests have their own conditions. These heterogeneous checks should not be compressed into a single universal residual.

The manifest records Python 3.12.14, NumPy 2.3.5, SciPy 1.17.0 and SymPy 1.14.0, the command, actual runtime, source hash, output hashes and the finite bounds. It validates as a legacy version 1 computation record. Version 1 validation does not by itself enforce machine-level resource caps or a complete dependency graph. This run adds its input hash and confirms it remained unchanged, but the interpretation stays “result reproduced; implementation not independently validated.”

12 The electrovac verification boundary

The charged spherical model defines the mass coordinate

$$\mathcal{M} = \frac{R}{2}\left[1 - \left(\frac{R'}{\Lambda}\right)^2 + \left(\frac{P_\Lambda}{R}\right)^2\right] + \frac{Q^2}{2R}.$$

Its exact identity expresses $\mathcal{M}'$ as a combination of the Hamiltonian, radial momentum and electric constraints. Recovering those constraints from the triangular mass-coordinate system needs nonvanishing coefficients on a monotone regular patch. The theorem is not an algebraic equivalence at every degenerate configuration.

The finite-range benchmark evolves a radius variable and reconstructs the other canonical fields from mass and charge. That projection can make reduced mass-shell residuals close to roundoff. Standard ADM residuals evaluated with independent finite-difference derivatives remain nonzero and converge. A roundoff-level projected residual must therefore not be advertised as exact agreement of all discrete ADM operators.

The retained benchmark uses mass 1, charge 0.8, a radial interval from 0.6 to 6, a smooth coordinate deformation, final time 0.4 and CFL factor 0.10. Its refinement grids have 64, 128, 256 and 512 nodes. The inherited ledger has 19 passes, including symbolic identities, fourth-order reconstruction fits, horizon-coordinate consistency, a stability diagnostic and an unprojected negative control.

The radius evolution stencil is finite range. A radial constraint prefix solve and the reconstructed geometry have a different locality property. The full calculation is therefore not a proof that every field update in the reduced gravity solver is strictly finite range.

The new bounded strong-field audit calls the source's symbolic electrovac identities and obtains exact zero. It also checks an independently written derivative identity needed by the regular source estimate. It reads the existing convergence and check data without rerunning the full evolution. The distinction is recorded in its manifest.

The uniform stability route still needs lower bounds for the radius derivative and the square-root radicand, controlled coefficients, boundary data and a suitable common norm. An observed small spread of amplification over the tested grids is useful evidence for that route. It is not a theorem uniform over all refinement levels or over nonspherical data.

13 The scalar-collapse evidence

The massless spherical scalar uses the flat-slice Painlevé–Gullstrand chart with $v = \sqrt{2m/r}$. Its exact reduced identities include

$$\partial_r m = r^2 \left[\frac{1}{2}(\Pi^2 + \Phi^2) + v\Pi\Phi\right],$$

$$\partial_t m = r^2\sigma(\Pi + v\Phi)(\Phi + v\Pi).$$

The symbolic compatibility calculation relates the time derivative of the radial mass constraint to the radial derivative of this flux. It uses the actual lapse constraint and scalar evolution on the regular region $r > 0$, $\sigma > 0$, $v > 0$. An extension to vacuum points $v = 0$ is separate. The continuum proof does not establish consistency of thresholded or clipped numerical reconstruction there. The outgoing null expansion vanishes when $2m/r = 1$ in the declared chart. This identifies a marginally trapped sphere, not a global event horizon.

The retained strong-packet table has seven grids from 600 to 2,400 nodes. Every listed strong run reaches the trapped-sphere criterion. A weak packet stays untrapped through time 5. The event-coordinate fits use the six finest grids, from 800 to 2,400, and give orders 2.173853 for time and 2.084395 for radius. The600-to800 event time is nonmonotone. The fitted evidence is therefore not an all-grid monotonicity or asymptotic-order theorem.

At the finest grid, the interpolated event values are

$$t_H = 4.6918209458, \qquad r_H = 3.1134115480, \qquad m_H = r_H/2.$$

The stored field diagnostics come from the later stopping time 4.6920696969. Its gap from the interpolated event is approximately 2.48751×10^{-4}. The lapse, curvature, outer mass and event-field profile should be identified with that stored field time unless explicitly interpolated. Combining all of them under one exact event-time label loses this distinction.

The finest reported relative residuals are approximately 2.2732×10^{-6} for the radial mass constraint, 2.8077×10^{-6} for the lapse constraint, and 1.2697×10^{-3} for the independent finite-time mass flux. The outer mass change is about 0.2624 percent. These are controlled numerical residuals, not exact conservation in the discretized evolution.

The spatial characteristic stencil has radius two. Geometry is reconstructed by triangular radial prefix solves at every RK stage. Those solves are not strictly finite-range maps. The benchmark has excision rather than regular-center reflection, and it does not establish a critical exponent, singularity continuation, nonspherical stability or a complete $3+1$ solver. The new symbolic audit reproduced the reduced identities but performed zero full collapse-evolution steps.

14 Reproducing the evidence without confusing the wrappers

The edition ZIP preserves the exact earlier archives in `recovered_sources/`. To inspect an inherited gravity claim, extract the Volume I source archive into a new working directory. Its root is `Projector_Refinement_Physics_Volume_I_Source/`. Use that extracted root as the working directory for its source commands. This avoids writing regenerated outputs over the preserved archive or over the edition evidence.

The technical supplement's labels V1E47 and V1E54–59 are code numbers, not chapter numbers. V1E54–56 reuse retained JSON by default. They invoke the underlying verifier only if its JSON is missing or the respective `RERUN_ASHTEKAR`, `RERUN_LORENTZIAN_BV` or `RERUN_HEAVY_BFV` flag is 1. V1E57–59 convert existing results into ledger rows. Re-executing a converter is not a new evolution or symbolic reproduction.

The inherited archive does not pin a Python dependency environment. Ashtekar verification uses JAX with its 64-bit option enabled; reduction uses NumPy, SciPy and SymPy; scalar evolution optionally uses Numba. Historical library versions are not recorded in the selected outputs. The current reproduction records its own versions and hashes.

The distinction is especially visible in V1E57. It reads `resolved_second_class_reduction_checks.csv` and selects existing pass rows. Several standardized residual fields are set to zero while the actual recorded value is retained in the note. For example, the RK4 drift is a successfully detected negative control, not a zero-drift numerical result. The raw source-key value is the meaningful quantity.

To bypass the three retained-JSON caches, run the underlying programs from the extracted root. The nonlinear transfer code and second-class verifier can also be run directly:

```
python code/47_nonlinear_bfv_bv_polyphase_bridge.py
python verification/Resolved_Ashtekar_Closure_Verification.py
python verification/Resolved_Lorentzian_BV_Parent_Verification.py
python verification/Resolved_Lorentzian_3plus1_BFV_Verification.py
python verification/Resolved_Second_Class_Reduction_Verification.py
```

The second-class program writes its JSON next to itself; the retained publication copy is in `data/`. Compare payloads rather than assuming identical locations. The accompanying record copies the unchanged verifier into a dedicated evidence directory and records its output there.

The underlying strong-field programs accept an explicit output directory:

```
python verification/Finite_Range_Spherical_Electrovac_Verification.py \
  --output /absolute/path/to/electrovac_run
python verification/Finite_Range_Spherical_Scalar_Collapse_Verification.py \
  --output /absolute/path/to/scalar_run
```

The scalar verifier also exposes run controls through its command-line parser. Changing grids, time bounds, packet amplitude or controls changes the tested assertion. Record them rather than comparing an altered run to the archived seven-grid result as if it were identical.

For the current edition's narrower checks, use its root directory:

```
python editorial/gravity/collect_verification_sources.py
python editorial/gravity/run_second_class_verification.py
python tools/verify_strong_field_identities.py
python tools/verify_gravity_brst_parity.py
```

The first command only reads ZIP members and regenerates the crosswalk. The other commands are bounded and record their own distinct claims. None invokes Lean.

15 The inherited Lean environment and its limits

The formal archive is `Komorebi_Physics_Lean_v0.49.0_Checked_Block_49.zip`. Its toolchain is `leanprover/lean4:v4.34.0-rc2`, and the Mathlib revision in `lakefile.toml` is

`aeb72f1e9981b6d35b426b16dca7759be3e7ba06`.

The retained package metadata still says version 0.48.0. That is a metadata lag in the Block49 archive, not evidence of a different theorem or a new compilation in this account. Both files are preserved and hashed.

The companion's source search covers 146 `KomorebiPhysics/*.lean` members in that archive. It searches source text for the direct gravity terms and examines the eight selected modules. It identifies no direct C08–C14 gravity declaration in that scope. This is a scoped source finding. It is not a semantic proof about every imported dependency, and it does not use the older global 142-module inventory as a current whole-project count.

The newest selected formal source, `Relativistic3DWeightedDiracDomain.lean`, is included to fix the actual frontier. In namespace `KomorebiPhysics.Relativistic3DWeightedDomain`, `mem_physicalDiracDomain_iff_summable_h1` identifies the maximal free physical domain with the first-order weighted Fourier condition. The declarations `physicalDiracPMap_selfAdjoint` and `physicalDiracPMap_isClosed` establish the self-adjoint and closed free operator. These are free torus statements. They do not include a dynamical tetrad, torsion, lapse, shift or gravity constraint surface.

The previous gauge companion already records retained-receipt exceptions: the Block 27 cumulative-build log contains a Git exit-code 128 failure inconsistent with the receipt, and six reported Blocks 36 and 40 logs are missing. This gravity edition preserves that finding and does not convert it into a mathematical counterexample. The direct source crosswalk here is independent of claiming that every inherited cumulative build receipt is complete.

A fresh formal reproduction would require the exact Lean and Mathlib environment and the archived release-specific commands. It must record successful source compilation and an axiom/trust audit. Merely running XeLaTeX, searching for a theorem name or observing a zero Python residual does not provide a new Lean check. No such execution was performed here.

16 What remains to prove

The closest new formal target is the finite discrete-gradient algebra. It would define exact discrete gradients, prove the endpoint tensor is skew and annihilates the constraint rows, and derive exact invariant preservation for a solved step. It should keep existence of the implicit solution separate. This target is narrow enough that its assumptions and output can be named without claiming a full gravity solver.

The graded parent needs a complete parity-consistent component audit. The newly detected Lorentz dual-spinor defect is repaired locally. The next check must cover the same ghost-placement convention in the diffeomorphism and electric sectors, the parent action, antifield derivatives, boundary moment maps and ghost momenta. A pure-Lorentz witness does not discharge those dependencies.

The matter realization also needs physical massive-Dirac convention matching. The inherited mostly-plus Clifford matrices with flat operators $i\gamma^I\partial_I \pm m$ give the wrong sign for the conventional massive shell. A repair must treat the kinetic factor, adjoint, mass convention and any induced torsion signs together; covariance alone cannot select it.

The resolved vacuum algebra needs an explicit formal Fourier-support calculus, canonical bracket, constraint densities, and a verified translation of the continuum self-dual identities. The support theorem must apply to all intermediate products and generated smearings. Uniform propagation of headroom is a separate dynamics problem.

The Lorentzian canonical chart needs a formal treatment of nondegeneracy, graded Dirac brackets, torsion elimination and boundary-charge tangency. Extending the chart across degenerate triads or constructing a global gauge quotient requires additional geometry. The local time-gauge formulas do not supply that extension.

The infinite tower needs level-independent regularity, coercivity and inverse-Dirac-matrix bounds; compatible evolution and reconstruction maps; and explicit summable defects in a common complete space. Only then can the already formal metric-limit theorem be instantiated for an Einstein–matter trajectory.

The spherical numerical route needs independent solver reproduction, a proof of stability and consistency for the complete projected system with its boundaries, and a quantified event-location limit. A global trapped-surface theorem would require additional continuum existence and geometric estimates. Regular-center evolution, critical collapse, singularity formation, event horizons and nonspherical stability remain further targets.

The quantum route needs a regulator, measure and quantum master equation with anomaly control. It is not completed by the classical transfer formula or the exact finite nonlinear example. Empirical validation and a comparison to independent physical data are further tasks beyond this mathematical verification companion.

17 A reader's verification contract

For any result in this gravity set, first identify the variables and the algebra. A Hermitian matrix derivation, a filtered Fourier field, a real tetrad, a half-density spinor and a spherical mass coordinate are different objects. Transfer between them requires a declared map and preserved structure.

Next read the domain. The essential gates include Fourier headroom, positive triad determinant, regular gauge fixing, endpoint-matrix invertibility, monotone radius, positive reconstruction radicand and continued membership in an excised chart. Some are pointwise conditions, some are neighborhood conditions, and some must be uniform over levels. The proof status changes when one of these conditions is removed.

Then identify the evidence type. For a symbolic zero, inspect the expression and arithmetic. For a floating residual, inspect the normalization, threshold, grid, seed and negative control. For an inherited pass row, inspect its method and whether the executable actually computed the claim. For a formal declaration, inspect the types and hypotheses rather than the physical phrase in its filename.

Finally trace dependencies. A classical master claim needs a consistent graded differential. A reduced master claim needs regular reduction and tangency. Exact invariant preservation needs an exactly solved step. A continuum claim needs uniform estimates and compatible reconstruction. A quantum claim needs its measure and regulator. The new BRST counterexample shows why those dependencies must remain explicit even when a large inherited ledger contains many passes.

The accompanying record provides a coherent gravity monograph, a Python visual atlas and this companion as three views of the same source. The manuscript gives the ordinary arguments, the atlas shows selected geometry and bounded numerical behavior, and the companion makes the remaining formal bridge and the evidence limits inspectable.

17.1 Portable evidence index

- `editorial/gravity/verification_source_map.json`: original member paths, hashes, seven claim groups and selected Lean locations.
- `evidence/gravity/second_class_computation_manifest.json`: fresh bounded underlying reduction reproduction.
- `evidence/gravity/strong_fields_identity_checks/`: symbolic strong-field identities and retained-table consistency checks.
- `evidence/gravity/brst_parity_results.json`: exact negative control and corrected pure-Lorentz dual action.
- `editorial/gravity/verification_layout_review.json`: final companion layout and proofreading scope.

The Adic Verification Companion

Exact sources, bounded checks, and the formal interfaces

1 The actual verification frontier

This companion records what has been proved, what has been computed, and what remains to connect the adic program to the formal physics archive. The finite construction uses ordinary complex amplitudes on finite address sets. Its exact statements concern unit coordinates, normalized embeddings, packet projectors, compatible Hamiltonians, Schur complements, and analytic residue lanes. The infinite construction has additional measure, domain and convergence hypotheses.

The adic source correspondence is limited to related modules in the preserved Block 49 archive. No Lean invocation was made for that adic audit, and its archived source was not changed. The separately checked free-Dirac generator does not certify the adic spectral program. Eight related modules supply normalized cell reconstruction, measure-preserving address coding, actual dyadic projections, aggregation composition, and a complete-metric refinement bridge. Their thirty-one selected declaration locations are recorded in the source map. They do not constitute a formalization of the complete adic spectral program.

The verification title is therefore deliberate. The monograph contains ordinary mathematical proofs, the atlas contains fresh bounded numerical checks, and this companion adds fourteen exact finite assertions. The inherited computational ledgers are retained separately. A successful identity check does not authenticate a global infinite-depth claim, and an archive hash does not authenticate a new kernel execution.

1.1 The distinction that controls the collection

A normalized embedding satisfies $I_h^* I_h = 1$. A compatible self-adjoint Hamiltonian also satisfies

$$H_{h+1} I_h = I_h H_h.$$

Only after both properties hold does the finite functional-calculus conclusion follow. A unitary fine walk alone does not supply this intertwining. In particular, extending the depth of the archived Grover product graph changes degrees and fails under the natural inclusion of its old arcs. Section 13 gives an exact witness. This failure concerns that inclusion, rather than unitarity at either fixed depth.

The portable crosswalk is `editorial/adic/verification_source_map.json`. It pins 116 selected archive members, including source chapters, executable programs, data, formal modules and receipt logs. Every member is identified by its original ZIP path and SHA-256 digest. These are evidence locators, not 116 independent mathematical results.

2 Claims and stable source locations

The technical supplement groups this material under C15 and C16. Its printed chapter ranges are stale relative to the current source package. The filenames below are stable locators. A chapter title or a compiled position must not replace an exact member path when checking the evidence.

ID	Source family	Exact scope of the current result
C15	Unit registers, chapters 29–31	Finite unit-coordinate bijections, CRT permutation unitaries and compatible multi-prime Hamiltonians. Complex Born amplitudes.
C15	Normalization and product walks, chapters 32–33	Counting-norm normalization and unitarity at each fixed product graph. Depth compatibility requires its own proof.
C15	Radial and polyphase reduction, chapters 34–35	Finite Jacobi and block reductions. Schur formulas hold where the eliminated resolvent exists.
C16	Infinite Haar limit, chapter 36	Dense cylinder union, strong projection convergence, and self-adjoint diagonal hierarchy with its actual domain.
C16	Prime decimation, chapter 37	Exact finite Schur composition on common regular charts, with a polynomial trace-coordinate composition law.
C16	Companion packet lanes, chapter 41	Entire residue-lane identities and bounded operator functions. Unbounded operator extensions need separate domain control.

Here “chapter 29” means the source filename beginning `29_`, rather than a claim about the current compiled chapter number. The full names are recorded in the machine-readable crosswalk. C15’s executable path is precisely V1E38–V1E44, seven programs. C16 selects V1E45, V1E46, V1E49 and V1E50, four programs. Together they contribute eleven retained ledgers.

The monograph also makes the complex-coupling convention explicit: a reverse transition must carry the complex conjugate coefficient for a Hermitian Hamiltonian. A real-coefficient example cannot verify an unrestricted complex-coefficient statement. Likewise a regular radial branch relation requires its energy to lie outside the eliminated tail spectrum.

3 What the inherited ledgers actually record

The eleven selected CSV ledgers contain 143 rows labelled PASS. The count is a historical source inventory. No original program was rerun to produce these rows for the documented audit. Exact copies are provided in `evidence/adic/inherited_ledgers/`, and the source map records their archive identities.

Code	Result family	Rows	Largest recorded residual
38	Unit coordinates and CRT	8	7.45×10^{-15}
39	Finite Haar refinement	8	1.10×10^{-14}
40	Finite-adelic dynamics	9	9.93×10^{-9}
41	Normalization weights	26	1.89×10^{-15}
42	Product graph walk	7	1.31×10^{-15}
43	Radial self-energy	9	2.07×10^{-14}
44	Polyphase refinement	25	1.27×10^{-14}
45	Infinite-limit surrogates	12	3.71×10^{-14}
46	Prime decimation	14	1.02×10^{-12}
49	Hypergeometric packets	17	3.91×10^{-16}
50	Transfer trace coordinates	8	7.77×10^{-12}

The residuals are heterogeneous. Some are matrix norms, others are coefficient differences, fitted-order discrepancies, or zero entries attached to an analytic explanation. Code 40's largest value is an order-fit discrepancy, rather than a failed exact commutator. Its PASS labels must be read with each row's method and note. Comparing all eleven maxima to one universal tolerance would have no mathematical meaning.

Code 38 checks a complete local coordinate table but samples the multiplication law. Code 43 evaluates the visible branch relation only for eigenvectors with visible weight greater than 0.08. Code 45 evaluates finite cylinders, not an infinite matrix on its domain. Those restrictions remain part of the retained evidence even when their CSV statement is shorter.

4 Executable routes and their limits

The selected original programs are executable Python scripts, rather than the cached verifier wrappers used elsewhere in the gravity source. Their top-level statements construct examples, write tables and render figures. Running one imports its numerical and plotting dependencies and may perform substantial matrix work. The accompanying preparation record did not execute them as a group.

For an original reproduction, first extract the unmodified Volume I source archive into a working directory. The program resolves its root from `__file__`, then writes to that root's `data/` and `figures/` directories. An example command, run from the extracted source root, is

```
python code/39_adic_haar_quantum_refinement.py
```

To reproduce the finite-cylinder experiment rather than claim an infinite execution, use

```
python code/45_infinite_adic_haar_quantum_limit.py
```

These are reproduction routes, not commands already run for the documented audit.

Several scripts build rows with a literal PASS string and apply a threshold only after exporting the CSV. Code 44 exits if the largest residual exceeds 2×10^{-10}; code 45 uses the same aggregate bound, and code 46 uses 2×10^{-9}. Code 49's row builder marks any finite residual PASS, while its final check requires seventeen rows and a maximum residual below 5×10^{-10}. Code 50 checks eight rows and the same bound. A prematurely written CSV can therefore retain PASS labels even if a later final guard fails.

Code 40 ends by printing its table and has no final aggregate rejection guard. Its exact algebra, perturbative bounds, finite exponential calculations and fitted orders need separate interpretations. A zero process exit alone would not establish every printed theorem.

The portable package preserves the original source bytes and outputs, while the new exact checker reports each implemented assertion explicitly. It exits nonzero if an assertion fails. Dependency versions and tolerance choices still belong to a reproduction record. The old archive does not supply a fully pinned environment for every original numerical script.

5 Formal normalized cell reconstruction

Block 35's `Relativistic3DL2CellReconstruction.lean` builds an actual map from a finite complex Hilbert state into L^2 using disjoint measurable cells of volume one. The declaration `UnitVolumePartition.stepIsometry` is a linear isometry, rather than an abstract assumption that interpolation preserves norms. Its adjoint averaging satisfies `averaging_step`: interpolation followed by averaging returns the original finite state.

Block 36 replaces unit cells by cells of a positive common volume v. Its normalization is

$$c_v = v^{-1/2}, \qquad c_v^2 v = 1.$$

The exact source theorem is `EqualVolumePartition.normalization_sq_mul_volume`. The declarations `inner_normalizedCellIndicator` and `stepIsometry` then establish the normalized inner product and the finite-to-continuum isometry.

This gives a real formal explanation of the normalization used in an equal-child adic refinement. If a parent is split into p children, counting-norm coefficients must be divided by $\sqrt{p}$. On the continuum side, normalized indicators acquire the reciprocal square root of cell measure. The two descriptions agree after the appropriate cell-volume identification. The current formal module instantiates an octree, whose cell volumes are 8^{-n} and side lengths 2^{-n}, and proves

$$8^{-n} = (2^{-n})^3.$$

The source does not define $\mathbb{Z}_p$, its unit group, or a prime-indexed Haar character basis in these selected modules. Generalizing the reusable normalization theorem to an actual p-adic measure model still requires the relevant cells, measure and disjointness facts. A similarity between 8-child refinement and p-child refinement is useful guidance, not a source declaration that identifies the spaces.

The density theorem `Octree.finitePrefixSubmodule_dense` proves finite-prefix approximation in the actual octree L^2 model. Its conclusion concerns approximation of states. It does not assert compatibility of an arbitrary hierarchical Hamiltonian with the projections.

6 Formal address coding and measure

Blocks 37 and 38 make the dyadic address-to-space bridge concrete. An octant address is converted into three binary streams, and binary expansion sends those streams into the unit cube. The source proves continuity and surjectivity of the coding map. It also proves a negative fact that must survive any interpretation: binary coding is not pointwise injective.

The theorem `binaryPoint_not_injective` uses the two familiar expansions of a dyadic rational. Measure theory handles this ambiguity by identifying functions almost everywhere. The later measure-preserving theorem is therefore stronger and more suitable for L^2 than an incorrect claim of a pointwise bijection.

In `Relativistic3DBinaryLebesgue.lean`, the binary distribution argument proves its actual cumulative distribution and then measure preservation. The selected declarations include `binaryCDF_eq`, `binaryPoint_measurePreserving`, and `cubeCode_measurePreserving`. Pullback consequently preserves the L^2 norm in the relevant measure model.

The correspondence is between a symbolic dyadic address space and Euclidean volume on the cube. The adic monograph's finite unit-coordinate CRT permutation is a different construction. It is a literal finite basis bijection, whereas binary Euclidean coding has pointwise ambiguity and is used through a measure-preserving pullback. Neither construction should inherit the other's hypotheses by analogy.

A p-adic physical interpretation would require another explicit bridge: an actual p-adic address measure, the intended observable algebra, and any proposed spacetime realization. The current monograph establishes its internal complex Hilbert dynamics without assuming that such a realization has been supplied.

7 Formal projections and their topology

Block 39 defines the adjoint of Euclidean-to-octree pullback and uses it to reconstruct actual Euclidean approximation spaces. The theorem `cubeAdjoint_pullback` gives the left-inverse property. The module defines orthogonal projections onto closed levels, proves contraction, proves nesting, and proves strong convergence to the identity.

For levels $a \leq b$, `euclideanProjection_of_le` has the form

$$P_b P_a f = P_a f.$$

The theorem `euclideanProjection_tendsto` states

$$\|P_n f - f\| \longrightarrow 0 \qquad \text{for each fixed } f \in L^2.$$

Block 40 then identifies those levels with the closed span of actual dyadic-box indicators. The theorem `prefixCell_ae_eq_preimage_dyadicBox` identifies prefix cells and box preimages almost everywhere. It avoids an invalid pointwise assertion at the boundary of binary expansions.

The declaration `euclideanLevel_eq_euclideanDyadicLevel` is a subspace identity. The theorem `euclideanDyadicProjection_tendsto` gives strong convergence for the box realization. Thus the formal bridge includes actual normalized interpolation, actual Euclidean cells, and convergence of the actual projections.

Strong convergence is the topology needed for state reconstruction. It is not operator-norm convergence. If an infinite-dimensional space has a nonzero detail vector above every cutoff, then

$$\|1 - P_n\| = 1 \qquad \text{for every } n.$$

That observation is fully compatible with $P_n f \to f$ for each fixed f. A bound on a chosen regularity class is a third statement and needs its own weighted norm assumption.

The adic infinite-limit proof uses the analogous cylinder-filtration mechanism in its explicitly defined space. The selected Lean modules establish the dyadic Euclidean mechanism. They do not yet instantiate the prime-indexed Hamiltonian, its resolvent, or the companion residue-lane functional calculus.

8 The formal exact-tower interface

The module `LocalDiscreteGlobalContinuum.lean` defines an exact refinement tower through its actual required data: state types, embeddings, level evolutions and an adjacent intertwining hypothesis. Its declaration `evolve_iterEmbed` propagates

$$U_{n+1}I_n = I_nU_n$$

through any finite number of embeddings. It does not construct the embeddings or prove adjacent compatibility for every concrete walk.

The same module defines a `SummableDefect`. A reconstruction sequence f_n in a complete metric space must satisfy

$$d(f_n, f_{n+1}) \le b_n, \qquad \sum_{n=0}^{\infty} b_n < \infty.$$

The declarations `exists_limit` and `dist_limit_le_tail` then prove convergence and the quantitative bound

$$d(f_n, f_\infty) \le \sum_{j=0}^{\infty} b_{n+j}.$$

The declared complete target and summable sequence are load-bearing assumptions. A small finite residual does not manufacture a summable defect sequence.

The theorem `globalEvolution_error_bound` applies the reconstruction mechanism to the coherently evolved tower. For the adic model, a concrete use would need the specified prime-indexed spaces, normalized embeddings, compatible evolution and a reconstruction satisfying the defect requirement. A diagonal hierarchy that stabilizes old Haar weights gives one ordinary mathematical route. A depth-dependent graph walk with changing local coins requires another route.

`RefinementTower.lean` supplies two further exact building blocks: `push_comp` proves composition of finite aggregation, while `aggregateCurrent_comp` proves composition of the matrix current transfer $F \mapsto RFR^T$. These aggregation identities should not be renamed as Schur-complement transitivity. Aggregation sums variables; Schur elimination returns the effect of an inverted hidden block. The monograph proves the latter separately on its regular spectral chart.

9 Inherited formal receipt identity

The accompanying record compared selected current archive bytes against retained source hashes. The seven selected modules already mapped in the gauge companion match their inherited identities. The eighth, `RefinementTower.lean`, is pinned to the current archive and read directly; this companion does not add a new historical receipt comparison for it.

The narrowed Block 35–40 receipt audit contains 37 matching log hashes, six unresolved log paths, and three available logs without an expected receipt digest. This count concerns evidence availability and identity. It is not a new compilation count.

Block	Retained evidence boundary
35	Source, API, axiom, build and negative-control identities are available in the selected receipt scope.
36	Three cited build or source-compile paths are missing. Available API, axiom and negative-control logs remain usable as inherited identity evidence.
37–39	Selected source and retained receipt-log identities match. Block 39 uses a concrete Euclidean projection module.
40	Direct compile, module build and cumulative build log paths are missing. Three negative-control logs are available but unpinned by the original receipt.

The formal archive retains Lean 4.34.0-rc2 and the recorded Mathlib revision for these releases. The corresponding reports declare their trust and axiom policies. The accompanying record leaves that inherited compiler and cache boundary unchanged. Its source map does not infer freshness from the phrase “checked archive.”

A useful future kernel check would instantiate a specific finite Haar embedding and its Hamiltonian compatibility first. It should run the actual declaration and API check with the recorded toolchain, produce fresh complete logs, and pin those logs to exact source bytes. Recompiling unrelated formal modules would not close that instantiation gap.

An unresolved log is reported as missing evidence, rather than a mathematical counterexample. Conversely, a matching empty successful compile log is an identity match to its historical receipt, rather than proof that a compile just occurred. These distinctions let the collection preserve progress without upgrading incomplete provenance.

10 The fourteen fresh exact assertions

The script `tools/verify_adic_identities.py` performs small deterministic checks in integers, rational numbers and exact symbolic rational functions. It does not execute any original archive program. Its final bounded run reports fourteen passed assertions and zero failed assertions. Three of the checks are named negative controls: their success means that a proposed incorrect simplification was detected.

No.	Implemented assertion	Tested range
1–2	Rational averaging algebra and detail ranks	$p = 2, 3$, depth 3, at most 27×27
3	Incidence Gram normalization	$p = 2, 3, 5$, coarse dimension 3
4–5	CRT dual phase and missing-dual negative control	Moduli $(2, 3), (3, 5), (4, 9)$, all 1557 basis pairs
6	Cyclic-shift incidence intertwining	$p = 2, 3, 5$, coarse dimension 3
7–9	Resolvent return, coupling degree and determinant factorization	One exact symbolic 3×3 family
10	Nested Schur transitivity	One exact rational 4×4 matrix
11	Grover depth-extension leakage witness	$p = 2, 3$, old degree 3
12	Factorial residue-lane coefficients	$2 \le N \le 7$, all j, $0 \le q \le 8$
13–14	Reducing packet and coupled negative control	Rational dimension 4, polynomial degree 3

The negative controls detect the incorrect CRT output permutation, nonzero Grover depth-extension leakage, and off-diagonal packet coupling. The checker reports fourteen assertion groups, not fourteen unrelated theorems.

The current result and manifest are in `evidence/adic/identity_checks_final/`. The version 2 manifest verifies the final script hash before and after execution, declared result hashes, exit status and effective resource caps. It was validated against the live project root. Correct interpretation still depends on the bounded contract and the displayed ordinary proofs.

11 The CRT phase check and normalization

For coprime a, b and $N = ab$, the input CRT map sends n to $(n \bmod a, n \bmod b)$. The Fourier output must instead use

$$k_a = kb^{-1} \pmod{a}, \qquad k_b = ka^{-1} \pmod{b}.$$

Indeed the product Fourier kernel has exponent

$$2\pi i \left(\frac{k_a n_a}{a} + \frac{k_b n_b}{b} \right).$$

After putting over denominator ab, its numerator is congruent to kn modulo ab. Thus it equals the N-point kernel. The exact checker tests this congruence for all basis pairs of three selected coprime modulus pairs. It also detects that reusing the input residue map for the output fails in this family.

The test uses integer modular arithmetic, rather than floating complex exponentials. It therefore separates a phase convention mistake from a tiny numerical Fourier residual. The general congruence argument supplies the proof for arbitrary coprime moduli; enumeration supplies a bounded implementation check.

The normalization check has a similarly exact surrogate. Let B be the coarse-to-fine incidence matrix with one parent assigned to each fine address and p children per parent. Then

$$B^* B = p\, 1, \qquad I = p^{-1/2} B, \qquad I^* I = 1.$$

The rational checker verifies $B^*B = p1$ directly and keeps the square-root normalization outside the integer matrix. Raw replication increases squared norm by p; it must not be used as a counting-norm isometry.

For a common finest depth H, the conditional averaging matrices are rational:

$$(E_r)_{xy} = p^{-(H-r)} \mathbf{1}_{x \equiv y \pmod{p^r}}.$$

The checks give $E_r E_s = E_{\min(r,s)}$ and the orthogonal differences $D_r = E_r - E_{r-1}$. Their traces recover the detail ranks. These exact formulas supply independent evidence alongside the atlas's floating Haar-basis computations.

12 Return terms and regular spectral charts

The exact symbolic family used by the checker is

$$H = \begin{pmatrix} 1 & \lambda & 0 \\ \lambda & 2 & 1 \\ 0 & 1 & 3 \end{pmatrix}.$$

Retain the first coordinate and eliminate the remaining two. Direct inversion of the detail block gives

$$\Sigma(z) = \lambda^2 \frac{z-3}{(z-2)(z-3)-1}, \qquad \left((z-H)^{-1}\right)_{00} = \frac{1}{z-1-\Sigma(z)}.$$

The resolvent formula assumes both the detail inverse and the full inverse exist. The determinant factorization can be cleared to a polynomial identity, but a cleared identity does not make every intermediate inverse defined at a pole.

The checker obtains the self-energy independently from the block inverse, verifies the displayed compression of the full inverse, and verifies determinant factorization. It also compares direct and staged elimination of a separate rational 4×4 matrix. That comparison tests finite Schur transitivity under invertibility, rather than an infinite operator limit.

The quadratic factor λ^2 has a precise perturbative interpretation. If $\lambda = p^{-w}$, the self-energy is of order p^{-2w} with a constant controlled by the detail resolvent. Letting z approach a detail pole can destroy a uniform bound. A fitted small-coupling slope away from the pole does not close that issue.

The packet negative control addresses a different possible shortcut. A block-diagonal rational 4×4 Hamiltonian preserves its rank-two reducing packet under a cubic polynomial function. Adding a symmetric visible-to-hidden entry produces nonzero leakage under the same polynomial. Residue-lane identities alone do not force a spatial or packet projector to reduce an arbitrary Hamiltonian. The commutation or intertwining hypothesis remains essential.

13 A concrete depth-compatibility failure

The product graph in the source combines a cycle with a rooted regular tree and uses the Grover coin at each vertex. At a fixed depth, the coin and arc-reversal shift are unitary. Extending the tree depth changes the degree of an old leaf. This prevents the natural inclusion of old arcs from being an exact refinement intertwiner.

At an old leaf away from exceptional cycle sizes, the old degree is $d = 3$: two cycle arcs and one arc to its parent. After extension, p child arcs are added. The new coin is

$$G_{d+p} = \frac{2}{d+p}\mathbf{1}\mathbf{1}^* - 1.$$

Apply it to one old outgoing arc basis vector. Every newly added outgoing child arc receives amplitude $2/(d+p)$. Therefore the squared norm in the new-arc sector is

$$\|QG_{d+p}e_j\|^2 = \frac{4p}{(d+p)^2} > 0.$$

The shift only permutes arcs and preserves this norm. For the checked cases,

$$p = 2: \quad \|QG_5e_j\|^2 = \frac{8}{25}, \qquad p = 3: \quad \|QG_6e_j\|^2 = \frac{1}{3}.$$

The coarse evolution followed by old-arc inclusion cannot contain these new arcs, so the intertwining defect is nonzero.

This is an exact finite counterexample to one proposed embedding route. It does not invalidate fixed-depth probability conservation, the graph-distance propagation cone, or the hierarchical Haar Hamiltonian whose old sectors are explicitly stabilized. Those are different results with different operators.

A compatible graph tower could use different embedding data, boundary treatment, local coins, or quantitative approximate compatibility. Any such route needs a construction and proof. The complete-metric exact-tower theorem cannot accept a failed compatibility hypothesis. The atlas's fixed-depth product-walk example should consequently be read as a unitary finite graph illustration, with no automatic continuum trajectory.

14 Infinite depth: the domain and topology gate

In the monograph's cylinder realization, let P_h project onto the level-h subspace, and let D_r be its orthogonal detail differences. Density of the cylinder union gives $P_h f \to f$ in L^2. A hierarchical Hamiltonian with real detail weights λ_r is an unbounded operator when those weights grow. Its definition includes

$$\mathcal{D}(H) = \left\{ f : \sum_r \lambda_r^2 \|D_r f\|^2 < \infty \right\}, \qquad Hf = \sum_r \lambda_r D_r f.$$

A state can belong to L^2 and fail this domain condition. Strong convergence of cylinder projections alone does not imply convergence of $HP_h f$.

For example, choose one unit vector e_r from each nonzero detail sector and coefficients $1/(r+1)$. Then $\sum_r (r+1)^{-2} < \infty$, so the state is square integrable. With weights $\lambda_r = r+1$, its graph-norm sum contains infinitely many terms equal to one. It lies outside $\mathcal{D}(H)$. The obstruction is an ordinary infinite-series argument, not a numerical cutoff test.

For a nonreal spectral parameter, a self-adjoint diagonal hierarchy has a bounded resolvent. Finite cutoff resolvents converge strongly on fixed states when the compatible hierarchy is used. On the omitted sectors the resolvent norm is controlled by

$$\sup_{r>h} |z - \lambda_r|^{-1}.$$

If the weights tend to infinity in absolute value, this tail tends to zero. If the weights stay bounded, that particular operator-norm tail conclusion need not hold. A finite-state tail estimate and an operator-norm tail estimate are separate statements.

The original code 45 studies finite cylinders and prescribed coefficients. Its small residuals are useful checks of a finite surrogate. The uniform infinite conclusion comes from the monograph's orthogonal-series proof with explicit domain, rather than from the largest tested depth. The selected Lean bridge formalizes related dyadic projections, not this prime-indexed spectral operator.

15 Companion lanes and bounded operator functions

For $N \geq 2$ and $0 \leq j < N$, the exponential residue lane is

$$E_{N,j}(z) = \sum_{q=0}^{\infty} \frac{z^{Nq+j}}{(Nq+j)!}.$$

The source's root-of-unity filter separates these lanes exactly, and their sum reconstructs e^z. The factorial coefficient factorization is

$$(Nq+j)! = j!\, N^{Nq} q! \prod_{\alpha \in \mathcal{A}_{N,j}} (\alpha)_q,$$

where $\mathcal{A}_{N,j}$ contains $(j+r)/N$, $1 \leq r \leq N$, with its single parameter equal to one removed. This is the coefficient identity behind the generalized hypergeometric expression.

The exact checker evaluates all 243 selected coefficient instances: $2 \leq N \leq 7$, every residue j, and $0 \leq q \leq 8$. It uses rational products, rather than floating special-function evaluations. The monograph supplies the general product argument. The atlas separately compares high-precision series, root sums and hypergeometric values for selected inputs.

For a bounded operator A, the entire series defining $E_{N,j}(A)$ converges in operator norm. The exponential reconstruction then holds in the bounded-operator algebra. If a projector P commutes with A, it commutes with every norm-convergent power series in A. This is the actual reason an established reducing packet survives the lane operation.

For an unbounded self-adjoint operator, e^{itH} is defined by spectral calculus for real time. A root filter containing factors such as $e^{\omega_N^{\ell} zH}$ can include exponentials with unbounded real growth. Their domains may differ, even when a sum has a formally attractive scalar identity. The bounded result must not be transferred to all states of an unbounded hierarchy without a common-domain or spectral-integrability theorem.

The finite cyclic companion matrix realizes an analytic residue system. Its word "companion" does not identify it with a physical quantum subsystem or an additional Hilbert-space degree of freedom by itself.

16 Reproduction and the next formal bridge

The new exact checker can be run from the bundle root with

```
python tools/verify_adic_identities.py --output \
  evidence/adic/identity_reproduction/results.json
```

This command computes the same finite assertions and writes a new result file. A result file alone is not a fresh provenance manifest. For a recorded run, use the computation-audit runner with `evidence/adic/identity_contract.json`, pin `tools/verify_adic_identities.py`, declare the new result path, and choose a fresh output directory. The retained final manifest records the actual argv, software versions, caps, logs and SHA-256 digests.

The final run used a 45-second wall limit, a 35-second CPU cap, a 1 GiB address-space cap, one effective CPU core, conventional one-thread numerical-library settings, and a 1 MiB log limit. It completed successfully and its version 2 manifest validated against the project. Conventional thread settings can be ignored by libraries; the recorded risk remains explicit.

An initial run failed while serializing SymPy integer metadata to JSON. It produced no scientific result file. The serializer was corrected to write ordinary integers. That failure log and its exact input snapshot are retained. A subsequent successful run was superseded after the self-energy check was strengthened to compare an independently computed detail-block inverse. The final live source is pinned by `identity_checks_final/manifest.json`; `identity_run_history.json` explains all three attempts.

The next formal milestone is specific: define one finite prime-address space, its counting inner product, normalized child embedding and detail projectors; prove the projector algebra and the compatible hierarchical Hamiltonian; then prove resolvent or unitary intertwining. A later phase can formalize a prime-indexed measure realization, the dense cylinder union, the unbounded operator domain, and the spectral limit. The fixed-depth Grover graph needs a separately chosen refinement route.

Four later physical questions remain outside that milestone: identifying a spacetime realization, deriving a coupled fiber/backreaction model, supplying continuum field equations for that model, and testing physical predictions. The accompanying record provides a carefully stated mathematical substrate and reusable formal interfaces. None of its evidence counts closes those additional implications.

APPENDIX A

Computational Evidence and Reproducibility

A.1 Source integrity and executable provenance

The release audit found two different kinds of evidence. First, the Volume II source manifest is internally intact. All 280 listed entries match their recorded SHA-256 digests. Second, the computational package does not rebuild in place without staging and inherited files.

Record	Count	Meaning
Volume II manifest entries verified	280	byte integrity
Frozen internal checks recorded as passing	203	archived implementation evidence
Byte-identical inherited data artifacts	38	cross-volume provenance
Byte-identical inherited figures	4	cross-volume provenance
Nested scripts requiring a staging layout	12	not runnable in place
Required but absent declared paths	4	three chapter records and one release manifest

The 203 passing checks were counted from ten archived verification summaries. They are frozen records. They were not all rerun from the Volume II package because the namespaced programs resolve their working root to the verification directory and then expect absent sibling data or parent inputs. Three records named by the scale-covariant chapter are absent from that package. Their ancestors exist in the first-volume source. The release documents also require a release manifest that is not present.

The Volume I technical supplement contains a visible status-key typo caused by a tab escape replacing the backslash in several `\textbf` commands. This bridge records the defect in an errata ledger. The frozen release is not silently overwritten.

> **Provenance boundary**
>
> Hash agreement proves that a file has not changed relative to its manifest. Byte-identical inheritance proves lineage. Neither proves that the physical assumptions are correct. A frozen pass count is not an independent rerun.

A.2 Scope of the executable record

The accompanying package reproduces the numerical statements associated with this model. It also performs a read-only integrity and lineage audit of the extracted parent sources. It does not claim to rerun every frozen parent computation.

The bridge verification has three classes of output:

1. analytic constants evaluated in double precision
2. resolved numerical identities and deliberately failing negative controls

3. deterministic audits of hashes, paths, and cross-volume byte lineage

The machine-readable claim registry gives every central statement a status and a validation route. The parameter registry separates structural choices, benchmark values, and derived quantities.

A.3 Clean-run commands

From the package root on a Unix-like system, create an isolated environment and run

```
python -m venv .venv
. .venv/bin/activate
python -m pip install -r requirements.txt
python code/verify_covariant_backreaction.py \
  --parent-json parent/Scale_Covariant_Dark_Sector_Verification.json \
  --output data --figures figures
python code/generate_latex_values.py \
  --data data --output generated_values.tex
python code/validate_release.py --root .
latexmk -pdf -interaction=nonstopmode -halt-on-error \
  -outdir=build \
  Projector_Refinement_Physics_Covariant_Backreaction_Bridge.tex
```

The parent JSON is included in the source package under `parent/`. The separate provenance audit requires extracted Volume I, Volume II, and Volume I technical-supplement source roots. Its command and expected directory arguments are documented in `REPRODUCIBILITY.md`.

On Windows PowerShell, `build.ps1` performs the bridge verification, macro generation, validation, and PDF build after the Python environment is installed.

A.4 Verification checks

The new suite records nineteen passing checks. Five of them are negative controls that pass only when the deliberately damaged or uncontrolled case is detected.

ID	Test	Result	Boundary
COV01	finite-level commutant dimension	pass	numerical linear algebra illustration
COV02	refinement coefficient nullity	pass	numerical recurrence illustration
COV03	FLRW exchange identity	pass	double-precision residual
COV04	omitted exchange source	pass	nonzero negative control
ADM01	alias-free scalar-fiber bracket	pass	residual 9.907025×10^{-16}
ADM02	violated spectral headroom	pass	residual 1.8994349731
ADM03	centered-difference derivative	pass	residual 0.0357793042
ADM04	omitted detail gradient	pass	residual 1.2401043846
UNI01	$\sigma = 1$ closed forms	pass	four analytic constants
UNI02	outer quadrature at $\sigma = 0.5$	pass	independent continuum moments

UNI03	outer quadrature at $\sigma = 1$	pass	independent continuum moments
UNI04	outer quadrature at $\sigma = 2$	pass	independent continuum moments
QUA01	96-node versus 128-node rule	pass	energy and pressure convergence
BKR01	specified benchmark energy overdraw	pass	uncontrolled production is detected
BKR02	parent overdraw reconstruction	pass	direct discrete moment check
BKR03	order-one threshold root	pass	$R_{\mathrm{disc}} = 1$ residual
BKR04	raw-continuum identity	pass	independent continuum coefficient
CPL01	nominal loop-coupling bound	pass	$g_{\mathrm{peak}} = 1$ boundary
CPL02	energy-control coupling relation	pass	linked λ scaling

"Pass" for a negative control means that the deliberately damaged model produced a residual above its declared lower bound. It does not mean that the damaged model is acceptable.

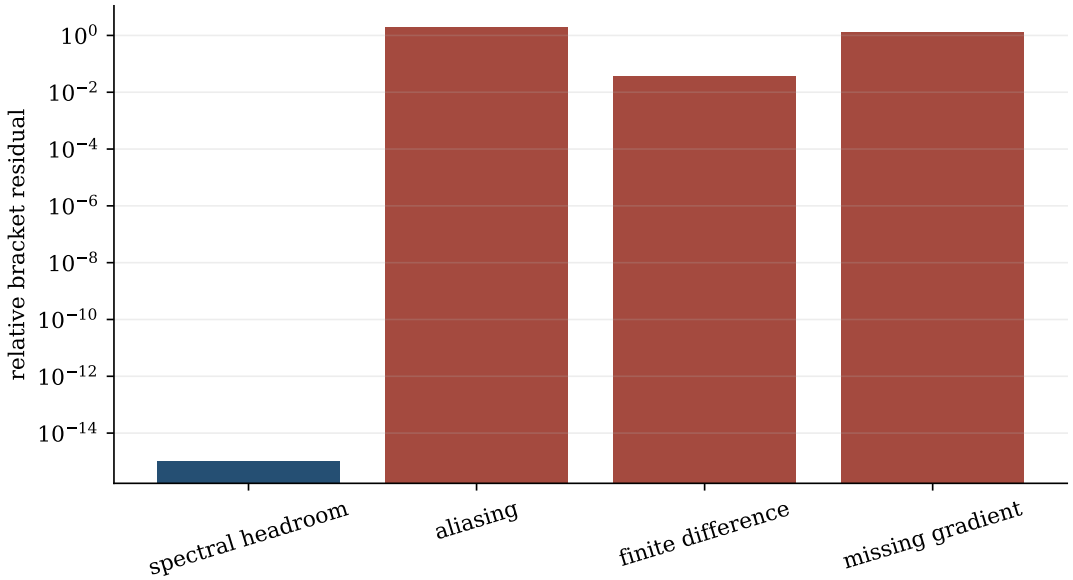


Figure A.1. Resolved bracket closure and negative controls. The alias-free spectral calculation closes near machine precision. Aliasing, a derivative without an exact product rule on the represented algebra, and removal of the detail gradient expose distinct failures.

A.5 Data products

Path	Contents
`data/verification_checks.csv`	individual check status, residual, tolerance, and role
`data/verification_summary.json`	aggregate pass count and central claim boundary
`data/backreaction_audit.json`	benchmark overdraw, order-one threshold diagnostics, and coupling bounds

`data/continuum_universality.csv`	arbitrary-σ analytic constants
`data/discrete_convergence.csv`	fixed-prime discrete and interpolation comparisons
`data/energy_saturation_curve.csv`	cutoff-controlled discrete energy, pressure, and rest moments along the diagnostic scan
`data/radiative_scaling.csv`	finite-depth ultraviolet moments and fitted scaling inputs
`data/frozen_verification_ledger.csv`	pass counts read from archived parent summaries
`data/cross_volume_lineage.csv`	byte-identical parent artifact pairs
`data/python_staging_audit.csv`	static root-resolution audit for namespaced programs
`data/errata_and_broken_references.csv`	missing declared paths and released erratum

All CSV files use a header row and UTF-8 encoding. JSON files are indented and key sorted. Figures are emitted as both PDF and PNG from the same Python calculation.

A.6 Source-archive audit

The parent audit verifies 280 Volume II manifest entries. It finds no digest mismatch. It counts 203 passing and zero failing checks across ten archived summary files. These are frozen results, not fresh executions.

The audit also detects 38 data files and 4 figures in Volume II that are byte-identical to Volume I ancestors. The result is positive provenance evidence. It is not independent confirmation because identical bytes carry the same assumptions and implementation history.

Twelve nested Python programs require a staged directory layout to resolve their declared root. Four declared release or chapter paths are absent. Those conditions are retained in the ledger rather than repaired invisibly.

A.7 Source erratum

The supporting technical supplement contains malformed status-label typography. The original file is retained unchanged in the provenance archive. The defect concerns presentation rather than a mathematical identity. Source identity and the corrected reading text are recorded separately.

A.8 Numerical methods and tolerances

Momentum moments use a 128-node generalized Gauss–Laguerre rule with weight exponent 1/2 after the production exponential is separated analytically. The pressure integrand supplies one additional factor of the quadrature variable. The discrete tower is truncated only after the dimensionless mass reaches a declared suppression cutoff. The order-one backreaction threshold is solved by a bracketed scalar root finder. Spectral derivatives use the periodic discrete Fourier transform.

Positive checks use tolerances appropriate to double-precision arithmetic or to the declared quadrature. Negative controls use lower bounds rather than upper tolerances. Every row states which direction is expected. The package validator rejects any failed row, a disagreement between the CSV and JSON counts, or a missing central output.

The computations are deterministic. The only randomized test uses a fixed seed, 20260819, and records the worst residual across its trials.

A.9 Reproduction boundary

Reproduction of the bridge establishes that the formulas were implemented as documented and that the reported diagnostics follow from the frozen input. It does not establish that the input action describes nature, that the chosen benchmark is unique, or that another implementation will agree outside the tested domain.

Independent scientific confirmation requires at least one implementation not derived from this code, a specialist check of the analytic derivations, and a new backreacting solver that does not impose the output spectral shape.

APPENDIX B

Claim and Parameter Ledger

B.1 Claim classes

The complete machine-readable ledger is `claim_registry.csv`. The principal classes are summarized here.

Class	Evidence required	Prohibited inference
Theorem	derivation from explicit assumptions	empirical truth without a physical identification
Conditional theorem	derivation plus declared regularity, cutoff, or resolution condition	validity after the condition is removed
Controlled diagnostic	deterministic result for a specified model and input	uniqueness or observational preference
Model proposition	internally consistent candidate with tests	derivation from projector axioms
Frozen record	verified archival bytes or archived check status	independent rerun
Open problem	explicit target and acceptance criteria	completed result

B.2 Parameter taxonomy

Category	Examples	Meaning
Structural	p, σ, H	selects the finite refinement geometry and spectral order
Microscopic	$\Lambda, m_b, v, \lambda, \gamma$	defines the effective action and interaction strength
State	initial $\phi, \dot{\phi}$, adiabatic vacuum order	defines a solution, not the theory alone
Cosmological	$H_0, \Omega_b, \Omega_D, n_s$	external background or fit input
Numerical	momentum cutoff, time step, quadrature order	must disappear from stable predictions
Derived	$A_\star, \eta, w, f_R$	follows only after the action, state, and approximation are fixed

The scale factor value $a_\star$ is not by itself an invariant physical parameter. A completed model should tie the crossing to a temperature, curvature scalar, Hubble rate, or another physical clock.

B.3 Integrity checks

The source integrity checks distinguish the following requirements:

1. all new verification rows pass
2. the parent manifest audit has zero hash failures
3. the claim and parameter registries contain no empty status field
4. every TeX input and figure dependency exists
5. the PDF builds without undefined references
6. every rendered page passes visual inspection for blank pages, clipping, malformed glyphs, and unreadable figures
7. the source archive expands into a clean directory and reproduces the same new data products within declared tolerances

These requirements concern artifact integrity, not scientific validation.

Bibliographic scope

The internal derivations are not an exhaustive specialist priority audit. Comparisons with hierarchical field theory, multiscalar effective models, preheating, scalar-tensor theories, constrained Hamiltonian gravity, and adic quantum models require their own source-level review. No literature-priority or empirical-preference claim is inferred from the breadth of this construction.

APPENDIX C

Reproducibility and Data Map

The source package contains the manuscript, figures, selected unified-theory data, and namespaced data and verification programs for the collisionless, CMB, thermal-history, lensing, nonlinear-structure, particle-mesh, halo-catalogue, baryonic-response, cooling-feedback, and subhalo-core calculations.

Build with `Build_Volume_II.ps1` or `Build_Volume_II.sh`. A clean rebuild and rendered-page comparison are included in the release audit. These establish package reproducibility, not external independent reproduction.

Bibliographic scope

These derivations and computations are not a complete specialist literature-priority audit. Hidden-sector towers, production mechanisms, scalar trapping, Boltzmann methods, CMB solvers, nonlinear structure, and galaxy-formation models each require their own comparative source review. The absence of an exhaustive priority audit limits historical novelty claims and is separate from the conditional mathematical results presented here.